



Telemetry-Driven Network Operations: Architectures for Analytics and Automated Remediation

Manevannan Ramasamy

Software Engineering Senior Technical Leader, Milpitas, California, United States of America

manevannan3@gmail.com

ABSTRACT: The network operation of telemetry is transforming conventional monitoring to diagnostic machine-aided, continuously, and automated remediation. The success of such operations, however, relies on an architecture that can handle heterogeneous telemetry at the correct time and space scales and whose operation is controlled in terms of collection, processing, storage, and operational costs. A telemetry-based network operations architecture presented in this study provides source-aware telemetry collection, event time standardization, topology-aware correlation, multilevel stream and past analysis, evidence-based diagnostic and policy-constrained remediation. The architecture is planned to organize the operational control loops by latency, confidence, and risk, such that a faster protective action can occur without the slower process of optimization and capacity planning. These implications enable an organized assessment framework to be defined using telemetry completeness, data freshness, and accuracy in anomaly detection, as well as latency in diagnosis, computational and storage overhead, elements of remediation accuracy, and policy adherence as significance measures in evaluation. The suggested method focuses on deciding the design of telemetry, where data gathering needs are based on operational purpose and not viewing telemetry as the end product. The architecture ensures accountable automation by having traceable relationships among observations, analytical results, choices, and actions, minimizing the chances of seizing the wrong remediation. As demonstrated in the paper, an integrated telemetry and analytics can be used to make the network more visible, responsive in its operations and remediation safety without sacrificing the explicit control as exercised by the human and policy authority.

KEYWORDS: Network Telemetry, Streaming Analytics, Automated Remediation, Observability, Event Processing, Network Operations

I. INTRODUCTION

Network monitoring history has substantially relied on an effort to periodically poll or check the device counters, interface statistics, logs and threshold based alarms. Such mechanisms are still helpful in making the daily supervision, but only they give a fragment of the picture of the network situation that rapidly varies. Fixed-interval sampling can ignore short-lived anomalies, and a solitary anomaly of the threshold can possibly lack sufficient context to determine whether an identified anomaly is a real degradation of service an intentional workload anomaly, or the realization of a deeper dependency failure. The modern network processes have consequently had to utilize constant, contextual and machine readable observations that can both carry out human analysis processes as well as the machine decision making processes.

Network telemetry Squid, Network telemetry can increase traditional monitoring with high-frequency streaming of operational information infrastructures and applications. Flow records, events, logs, traces, configuration state, routing data, topology relationships, service level indications can also feature as components of a telemetry environment in addition to classic counters. RFC 9232 provides a network telemetry model containing aspects of data production, export, gathering, analysis, and implementation to transform telemetry into an architectural attribute rather than a technique of just seeing statistics of devices [1]. Similarly, Openconfig gNMI too, employs long-lived subscriptions in which a network target updates in response to subscription modes such that a consumer can be informed of changes in operation without necessarily having to repeat polling the target only [2]. These functions provide ease of observability, but also present issues related to data volume, collection overhead, temporal alignment, storage, processing and operational governance.



The question of how to accumulate more telemetry is therefore not a key concern, but what becomes required is turning telemetry to valid operational evidence. A network would generate millions of measurements, events and records but would not be able to provide appropriate information to a particular operational decision. An architecture created based on data that has been over collected without an overall understanding of the operational objectives could end up being very expensive because of imposing additional bandwidth, processing and storage loads without a significant increase in the diagnostic. Conversely, excessive filtering, aggregation, or sampling could eradicate the contextual evidence that is required to identify infrequent failures, or between correlated symptoms and causes, or to confirm the effect of remediation. Effective telemetry architecture should thus establish a clear relationship between the questions that should be answered, the signals that should be received, the acceptable freshness, the analytics, and action.

What makes this requirement more important is that the operation of the network being automated in remediation. This automation process changes telemetry into a more service of observation, but not a source to provide control decisions. A latency jump can be sensed by a system of aberration-detection but an automatic response would require more support to determine that the situation is actable. Topology information may be employed to comprehend foreign dependencies; flow metrics might be employed to comprehend foreign traffic; configuration condition may be employed to comprehend that a change has just happened; and historical information may be employed to comprehend that the condition may be non-standard in the present workload. When there is no such contextual correlation, an automation system could be reacting to symptoms instead of yelling at causes, and this can further stabilize. Automation based on telemetry should thus create adequate evidence prior to authorization of an action and maintain the trace-ability between observations, analytical results, decision-making and changes to be implemented.

The second architectural concern is the temporal consistency. The telemetry sources could have varying frequencies of measurements and timestamps, buffering and collection characteristics. A network interface counter, an application-trace, a routing event and service-level indicator may indicate the same incident at various times. Event-time and correlation normalization are of importance in reconstruction of the sequence of the operations therefore. Lack of appropriate temporal processing may lead an analytics engine to a mistaken relationship where unrelated observations are wrongly linked, or a causal relationship, which occurs spread across many sets of data, gets mistaken. The architecture must also take into account the conditions of lack of telemetry, delays, duplications or a partial receipt since it is the quality of the data that will define the confidence level of analytical decisions.

Network size also makes telemetry processing a complicated issue. Big infrastructures can have thousands of devices, interfaces, services and logical dependencies and provide significant streams of heterogeneous observations. The needless processing time and resource usage may be added to the processing of each signal in single centralized analytic lines. A more appropriate architecture has the capability to spread analytics up and down many levels based on urgency and complexity of analytics. Regional or centralized processing can be used to perform topology correlation, historical analysis, capacity assessment and longer-horizon optimization; nearest-real-time processing can be used to identify anomalies and safety conditions on lower-horizons. This sort of separation allows a rapid control loop to respond, but does not require that all the analyses be placed in the same latency-critical pipeline.

The detection, diagnosis, decision and remediation should also be separated in the architecture. Detection is the point we identify deviation behavior of any observed behavior with any expected state. Diagnosis explores dependencies between signals in order to determine potential causes and dependencies affected. The intervention is defined depending on the existing policies to determine the intervention steps. Remediation does a proven activity, and thereafter, validates the outcome. Separating these steps into independent functions of architecture, increases a sense of responsibility and enables quantifying the location of mistake, or time waste in actions. It also facilitates different forms of autonomy to be attained depending on action risk and confidence of evidence and criticality of the service.

Using this, the paper describes a network operations architecture based on source-aware collections and event-time normalization, topology-sensitive correlation, tiered analytics, evidence-based discovery, and automating remediation using policy. The architecture considers telemetry as an operation evidence system of decision-expressed, rather than as an independent layer of data-acquisition. It is measured based on the criteria of telemetry completeness, telemetry fresh, accuracy of detection, diagnosis time, remediation safety and conformity to policy. The proposed method of linking telemetry demands to the operational goals and the explicit management of the shift between observation and action is aimed to enhance the visibility and responsivity of the networks besides ensuring the protection against unsuitable automation.



The rest of the paper will look into the architectural elements and how they affect one another, analytical and remediation processes and the criteria that have been utilized to evaluate the resultant operational structure. There is specific emphasis on observability efficiency, fast responsiveness/analytical trust and more with clear operational control. This view situates telemetry as a base to responsible network activity, as opposed to an augmenting, yet unregulated, mass of monitoring information.

II. RELATED WORK

Network telemetry has evolved to be one that does not involve a routine check-in but rather continuous data collection and analysis on operational data. A holistic network telemetry framework designed by Song et al. identifies its scope and includes the data generation, data export, data collection, data analysis, and data application that is the architectural foundation of the telemetry driven network operations [1]. The framework emphasizes the need to interlink the sources of telemetry and the analytical and operational processes instead of considering measurement collection as a standalone process.

To improve the streaming network management, the gRPC Network Management Interface (gNMI) has been added which offers configuration and operational data of the network devices by delivering them through subscription [2]. Subscriptions of this nature can save consumers the hassle of being updated and save on polling every time. However, to power up on streaming data, certain normalization, correlation, prioritization, and storage management mechanisms continue to be provided.

Besides network-related telemetry, OpenTelemetry provides a bigger model of observability, such as signals such as metrics, logs, and traces [3]. This multi signal viewpoint allows it to be operationally analysed both in infrastructure and applications, which is applicable in network settings where behavior of a service relies on distributed components. It is the heterogeneous signals that, however, make additional requirements of temporal and contextual, evidence management.

The other important trend is in-band Network Telemetry (INT) that involves integration of measurement information in network data paths. Survey carried out by Tan et al. takes into account INT mechanisms and how to exploit it to achieve fine-grained network visibility [4]. These methods can give detailed path-level information, although their implementation needs to take into account the overhead of telemetry, scaling, and the processing needs of high-volume measurements.

Automation has also been taken care of by model-driven management. The model, that Wu et al. propose is an automation of the service and network management by use of the YANG models, and instead preoccupies itself with standardized data model and automated processes in the field of management [5]. As demonstrated in this paper, domain frameworks can be found useful in aligning the network status with the management functions.

Network monitoring has significantly developed to become an alternative to traditional polling and rule-driven observation into telemetry driven, distributed and intelligent operational modes. Initial efforts put emphasis on the move towards scalable monitoring systems with the capacity to gather various forms of network data and assist with automated analysis [6]. State-transition detection then contributed to a system of finding important changes in network-element behavior allowing operating systems to react to the changing network conditions [7]. Both information about the control-plane and that about the data-plane has also been used to detect anomalies using telemetry, which shows the usefulness of sparse telemetry features to detect abnormal BGP behavior [8]. In networkwide traffic measurements, statistical methods were set to diagnose anomalies behavior in distributed network scenarios [9], and distribution mining of traffic features provided further evidence of the capability of detecting features characteristics deviation to show unknown abnormal patterns in the past [10]. In high dimensional telemetry streams, dimensionality is lowered using compressed PCA-based techniques keeping the important structures that detect anomalies [11]. Robust subspace learning builds upon this line with robust methods of PCA, subspace tracking, and recovery [12]. Covariance modeling has also been used statistically to characterize network traffic, and detect abnormalities which relate to flooding [13]. Together with these works, they put significant backgrounds in telemetry collection, state monitoring, feature extraction, dimensionality reduction, and anomaly detection. They, however, tend to focus more on personal tracking or analytical phases over combining telemetry ingestion with contextual analytics with decision-making and automated remediation to a single closed-loop network operation architecture.



The works combine useful foundations of telemetry gatherings, streaming, observability, meticulous measuring and automation. However, the synthesis of these capabilities in one architecture that explicitly links the quality of telemetry, smoke of analytics and the assurance of decision/policy-directed remediation is a major architectural challenge that this paper can speak about.

III. EVIDENCE-AWARE TELEMETRY FABRIC FOR ADAPTIVE NETWORK OPERATIONS

The proposed AI-minded telemetry fabric combines heterogeneous network sources and analytics, policy-informed policy and remediation operational findings. Rather than defining a telemetry collection as a rigorous line, collection fidelity, analytical latency and correction authority are persistently restructured to operational objective. It comprises of six major layers and they include source, transport, normalization, processing, evidence and remediation. Operation of freshness, completeness, confidence and ownership of the system would be provided in these layers through a dynamic telemetry system, which operates with clear operation semantics.




Figure 1 — Telemetry-Driven Network Operations Architecture

3.1 Source and Capability Layer

The raw observations made are provided in source layer needed to examine the network and services. It includes state of devices and interface statistics, routing tables, flow traces, system logs, synthetic tests, controller activities and application simplexes. There are wide differences in the sources in terms of sampling rates, nature of data, cost and reliability of data collection. This leads to each source publishing metadata regarding the capabilities of their sources provided; collection process accommodated, and expected event rate as well as resource requirements.

Subscriptions to telemetry are chosen based on the objective of the operation as opposed to just allowing as much fidelity as possible. As an example, wide-area infrastructure monitoring may require low-frequency interface and health statistics and active routing anomaly may require finer-grained routing and path data. The collection controller then considers, the diagnostic objective, expected rate of data, device capacity, what has already been done and what processing resources there are and then forms or modifies subscriptions. This technique enables the architecture to have high coverage but leave high-fidelity collection to be used in instances where there is more evidence that is operationally valuable.

3.2 Reliable Telemetry Transport Layer

The transport layer presents the schemes that guarantee the transportation of the telemetry in a reliable way, over the distributed sources, to the analytical services. It employs buffering, partitioning, compression, flow control and back-



pressure. These functions not only guarantee that slowdowns of processing transiently do not immediately propagate to network devices but also provide bursts which can be addressed in a managed way.

Provenance information, including the source identity, sequence, event time, collection time and schema version is provided in any telemetry event. Event time is used to determine when an observation was recorded and collection time is the time when platform has been sent or recorded an observation. This will be a big difference in the recreation of the sequences of the operations and the indication of the delayed telemetry. The transport layer then comes into play and distinguishes between observations which are delivered later and when the state is in fact new.

Duplicate processing matches as far as possible with stable event or sample identities. Rather than silently mismatch inconsistencies, the architecture forwards the loss indicators and data-quality metadata to the downstream processing. Components of analysis can, therefore, be used to establish a difference between lack of evidence and evidence that a condition had not taken place. This is a rule that does not allow incomplete telemetry to be viewed as a good condition.

3.3 Canonical Normalization and Context Layer

Normalization layer is an operation to convert vendor and protocol specific representations to canonical observation. Conflicts in items, counter semantics, timestamps, timestamps, sampling times, naming, and values are effectively sorted out before analytical comparison. This provides the down-stream correlation with a regular representation and preserves the record originals in case of forensic scrutiny as part of given retention policies.

Entity resolution connects telemetry observations that are time-versioned with time-versioned inventory and topology data. The fixed mappings cannot be helpful in historical analysis since the value of a device identifier, interface, service endpoint or logical network component may change over time. The architecture thus relates observations to the right timed variant of the infrastructure model.

3.4 Multi-Tier Analytical Processing

The level of processing separates the analytics according to the latency, the complexity of calculations and the risk of operation. It is suggested that there are three processing levels.




Figure 2 — Three-Tier Telemetry Analytics Model



Fast processing tier: This tier has latency at the seconds range and concentrates on instant protection signals. It performs lightweight anomaly testing, threshold testing, and health testing, and testing in a safety-oriented manner. The computational complexity and the set of actions are restricted on purpose because of the option to make a decision within a short period of time.

Operational processing tier: This is working within minutes and has more rigorous correlation in events, baseline, topology, dependencies and service indicators. It helps in diagnosing, recreation of incidences and controlled remediation. It can also base its work on the bigger contextual evidence to make suggestions on the change in its operations in comparison with the fast tier.

Planning processing tier: It is the processing tier whereby it evaluates the observations in hour, day or week. It defines general failure modes, capacity, general decreasing performance which is systematic and optimization of policies. Its outputs have more applications in assisting to make long-horizon operational decisions rather than immediate intervention.

Isolating these layers enables an efficient long term processing load to not interfere with safety critical processing and a clear resource budget and operational target in every control loop.

3.5 Evidence-Centered Operational Layer

The evidence layer converts correlated observations to action outcomes. The architecture uses temporal relationships, entities affected and services that are dependent on each other, and topology, to relate and cluster related observations as opposed to considering them as isolated events.

In each finding, there is a document of the compromised service or service purpose, date of incident, gravity, probable cause, quality of evidence and a candidate response. Topological relationships assist in deciding between an interface failure as a primary failure or just a downstream symptom. Temporal ordering also helps to differentiate initiating events and the one that follows.

Monitoring of alerts is therefore carried out at the finding and dependency level and not merely wipe down particular alerts. Several symptoms that a certain event may be connected with can be presented as one finding and the proves can be retained. This reduces repetitions of alert warnings without looking at the diagnostic environment.




Figure 3 — Evidence-Centered Diagnosis and Remediation Workflow



3.6 Policy-Governed Remediation Layer

Static full-fidelity collection can impose unnecessary bandwidth, storage, and processing costs. The architecture therefore supports three collection modes: **baseline**, **diagnostic**, and **forensic**. Baseline mode provides broad, low-cost coverage. Diagnostic mode temporarily increases sampling resolution around an active anomaly. Forensic mode preserves higher-fidelity raw evidence for detailed investigation.

Transitions between modes are themselves controlled actions with explicit rate, duration, and resource limits. Importantly, adaptive collection must not create circular evidence. When an anomaly causes sampling frequency to increase, the analytical system records this change so that subsequent increases in event volume are not incorrectly interpreted as worsening network conditions. Fairness controls additionally prevent a single noisy site or device from consuming the entire diagnostic capacity

3.7 Adaptive Telemetry Control

Full fidelity collection on rest can potentially be unwarranted bandwidth, storage and processing costs. The architecture thus facilitates three collection modes which are: baseline, diagnostic and forensic. Low cost coverage: It is comprehensive and is also available at a low price. Diagnostic mode temporarily increases sampling resolution in an area of localized anomaly activity. Evidence is stored at a higher fidelity in forensic mode to be examined in greater detail.

Mode transitions are explicitly-duration-limited also control actions. Interestingly, adaptive collection is not to produce circular evidence. In scenarios where the sampling rate increases because of an anomaly the analytical system informs later of increased volumes of this change, avoiding a false warning that network conditions are getting worse. Fairness controls also avoid one loud site or device crowding the available diagnostic bandwidth




Figure 4— Adaptive Telemetry Control Loop

3.8 Operational Semantics and Dual Observability Planes

Flow telemetry is new as far as the choice is concerned. An inventory record at the age of five minutes might not be necessary to implement a change of path but will be acceptable in reporting. Every analytical consumer thus stipulates a limit of acceptable age of data and limit of completeness. In case these conditions are not established, the resultant state is said to be unknown/less confident, and not assumed to be as healthy.

Platform observability is also separated by the architecture as compared to network telemetry. Platform observability is services that generate, collect, and process telemetry, whereas network telemetry is dictated infrastructure, traffic activities, and user experiences. OpenTelemetry has standard representations of distributed observability alerts such as



metrics, logs and traces [3]. This two plane map comes in handy, to find an apparent network failure in the network or in a collector, analytics service or a controller.

IV. ARCHITECTURE EVALUATION AND OPERATIONAL VALIDATION

4.1 Evaluation Dimensions

The dimensions that can be used to evaluate the proposed telemetry architecture to both analytical effectiveness and operational safety can be discussed. Telemetry completeness the portion of required observations that a decision can actually get, and freshness the propensity of observation timeliness within their age brackets to consumers. The significance of such metrics is that missing evidence or outdated evidence must make an analytical confidence lower instead of it being automatically considered healthy.

Detection accuracy determines how the architecture can detect relevant anomalies and minimize false notifications. The Latency of diagnosis is employed to show duration of time taken in making a transition between an initial abnormal observation on the one hand, and a well-founded operational result on the other. Resource efficiency is where the telemetry bandwidth, processing utilization, storage consumption and overhead brought about by adaptive collection are taken into account.

4.2 Analytical and Correlation Performance

It should assess whether using multi-source correlation is more helpful in yielding results than used on isolated signal processing. The role played by metrics, events, flows, topology information and service indicators may be taken into consideration, in order to determine which services are affected and also differentiating probable causes and downstream symptoms.

Special focus ought to be directed towards temporal correlation. The architecture is required to adequately react to the delayed, duplicates or reorder without generating conflicting results. It is also necessary to evaluate the baseline and diagnostic collection modalities to see whether telemetry high resolution will help increase the quality of evidence without causing disproportional resource utilizations.

4.3 Remediation Safety

Automated remediation needs to be assessed besides, detection. Under each candidate action, evaluation of the authorization compliance, confidence and blast radius, post-action verification and rollback ought to be performed. The assessment needs to assess the ratio of undertaken actions that are adequately evidenced and the rate of unwarranted or redundant interventions.

Actions that change configuration should be contrasted with least invasive ones e.g., additional diagnostic collection or path-testing. This distinction enables operational risk to be examined based on the extent of control that automation had.

4.4 Adaptive Collection Efficiency

Adaptive telemetry mechanism as a comparison of baseline, diagnostic and forensic mode should be taken into account. Indicators of interest will be a reduction in the average volume of telemetry, pinpointing during cases, signs of acquired diagnostics and utilization of resources per discovery.

The other measure of fairness is to determine whether diagnostic resources in the monitored sites remain the same and they are not monopolized by noisy incessant sources. This system must also make sure that transitioning between collection-modes is recorded, and does not appear to naive understanding of alterations to telemetry-volumes.

4.5 End-to-End Operational Assessment

End-to-end validation of typical operational cases in an early telemetry generation should be implemented through normalization, correlation, finding creation, policy evaluation, remediation and outcome verification. This cannot simply be about the occurrence or non-occurrence of anomaly but rather the entire chain must make a difference in prompt, telling and policy adhering response.



Combined the dimensions offer a balanced evaluation of visibility, analytic quality, efficiency, responsiveness and remediation safety such that the architecture can be viewed as an operational control system but not just an operational telemetry-gathering platform.

V. DISCUSSION

Network operation Telemetry Telemetry-based network operation is most effective when operational decisions are made and used to derive telemetry requirements (instead of the capabilities of the individual data sources). Action-oriented design is triggered by setting the action or the supposed question of decision, the evidence that is required to act as the decision and therefore, appropriate sources of telemetry, sample rates, retention and analysis methods are chosen. The existence of such a procedure will prevent the accumulation of large quantities of data blindly, simply because it can be created by network devices.

The scales of the telemetry resolution required by various operational options are different. Aggregated utilization measures in rather long periods can be good enough to make a decision about capacity-planning. On the other hand, an intermittent service disruption can be diagnosed by per-client, per-flow, per-path, or high-frequency measurements. Similarly, current interface, routing and topology information may be required by a fast remediation but historical trend analysis can endure the old information. Thus, telemetry freshness and granularity needs to be determined based on the latency and risk mandates of the decision.

The design also shows that quality of telemetry is an essential factor to be taken into account as a part of operational reasoning. The missing, late or absent observations cannot just disappear into an analytical pipeline. Instead, they ought to impact confidence and where proper, make the resulting state become unknown. This distinction is very important particularly to automated remediation since absence of evidence cannot be construed to mean a normal functioning.

Adaptive collection also provides a mechanism of a visibility vs. resource tradeoff. The broad baseline telemetry would be able to be employed to keep the coverage intact at all the time; diagnostic and forensic would be temporarily applicable when additional evidence is justified. However, this sort of transition should be controlled to avoid any kind of feedback between the outcomes and the inputs between some form of increase of collection also emerges like an oddity in itself running of operations.

Lastly, dashboards must be considered consumers of telemetry and not as the main motivation to collect information. Infinite high fidelity collection to populate already exuberant dashboards can provide access to large network, processing and storage capacity that has not resulted in improved operational outcomes. The telemetry design around explicit actions, evidence requirements and quantifiable decisions, therefore, bring a more long-term foundation of network operations by automation.

VI. CONCLUSION

Adaptive telemetry is another approach for balancing observability and resource efficiency. In situations where additional evidence is justified, a wide coverage is achieved by baseline collection, and the diagnostic and forensic modes are selective to narrow the data to a finer degree. Adaptive collection lacks the possibility of fairness limitations and guided transitions that would provide false feedback or take control over available resources.

This is further divided between automated action and the results of analysis regarding the architecture. Evidence-based discoveries sum up paralleled observing and policy regulation verification to approve, trust, fulfill a deadline, blast radius, rollback claim, and file validation requirement before remediation. The approach complies with automation and warns against operating responsibility.

In general, the idea of telemetry is most beneficial when gathering is geared toward clear operational choices as opposed to the dashboard or unlimited building of data. Future areas of work can delve deeper into the architecture by learning-based anomaly detection, more advanced models of confidence, auto-optimization of the policy, and standard evaluation of remediation results in a wide variety of network environments



REFERENCES

- [1] H. Song *et al.*, “Network Telemetry Framework,” RFC 9232, May 2022. [Online]. Available: <https://doi.org/10.17487/RFC9232>
- [2] OpenConfig, “gRPC Network Management Interface Specification.” [Online]. Available: <https://openconfig.net/docs/gnmi/gnmi-specification/>
- [3] OpenTelemetry, “Signals.” [Online]. Available: <https://opentelemetry.io/docs/concepts/signals/>
- [4] L. Tan *et al.*, “In-band network telemetry: A survey,” *Computer Networks*, vol. 186, Art. no. 107763, 2021. [Online]. Available: <https://doi.org/10.1016/j.comnet.2020.107763>
- [5] Q. Wu *et al.*, “A framework for automating service and network management with YANG,” RFC 8969, Jan. 2021. [Online]. Available: <https://doi.org/10.17487/RFC8969>
- [6] S. Lee *et al.*, “Network monitoring: Present and future,” *Computer Networks*, 2014.
- [7] P. Foroughi *et al.*, “Destin: Detecting state transitions in network elements.”
- [8] J. Cordova-Garcia, “Sparse control and data plane telemetry features for BGP anomaly detection,” in *Proc. IEEE INFOCOM*, 2019.
- [9] A. Lakhina *et al.*, “Diagnosing network-wide traffic anomalies,” *ACM SIGCOMM Comput. Commun. Rev.*, 2004.
- [10] A. Lakhina *et al.*, “Mining anomalies using traffic feature distributions,” *ACM SIGCOMM Comput. Commun. Rev.*, 2005.
- [11] Q. Ding *et al.*, “A compressed PCA subspace method for anomaly detection in high-dimensional data,” *IEEE Trans. Inf. Theory*, 2013.
- [12] N. Vaswani *et al.*, “Robust subspace learning: Robust PCA, robust subspace tracking, and robust subspace recovery,” *IEEE Signal Process. Mag.*, 2018.
- [13] D. S. Yeung *et al.*, “Covariance-matrix modeling and detecting various flooding attacks,” *IEEE Trans. Syst., Man, Cybern. A*, 2007.