



PDF Download
3737898.3769040.pdf
13 February 2026
Total Citations: 0
Total Downloads: 79

Latest updates: <https://dl.acm.org/doi/10.1145/3737898.3769040>

Published: 04 November 2025

[Citation in BibTeX format](#)

ACM MobiCom '25: The 31st Annual
International Conference on Mobile
Computing and Networking
November 4 - 8, 2025
Hong Kong, China

Conference Sponsors:
SIGMOBILE

SHORT-PAPER

PHOENIX: Protocol State Machine Recovery from Events in 5G Logs

QINGYUAN ZHENG, University of California, San Diego, San Diego, CA, United States

USHASI GHOSH, University of California, San Diego, San Diego, CA, United States

ALI MAMAGHANI, University of California, San Diego, San Diego, CA, United States

SRINIVAS SHAKKOTTAI, Texas A&M University, College Station, TX, United States

DINESH BHARADIA, University of California, San Diego, San Diego, CA, United States

Open Access Support provided by:

University of California, San Diego

Texas A&M University

PHOENIX: Protocol State Machine Recovery from Events in 5G Logs

Qingyuan Zheng¹, Ushasi Ghosh¹, Ali Mamaghani¹, Srinivas Shakkottai², and Dinesh Bharadia¹
¹University of California San Diego, CA, USA, ²Texas A&M University, TX, USA

Abstract

Modern multi-vendor cellular systems generate massive volumes of debug logs, while limited source code access makes protocol behavior modeling and fault diagnosis challenging. Traditional KPI-based analysis fails to capture protocol semantics, and existing log-driven inference approaches either suffer from high computational overhead or rely on oversimplified prior knowledge. We propose PHOENIX, an automated framework that reconstructs component-level finite state machines (FSMs) directly from raw network logs. PHOENIX constructs causal graphs from event sequences, applies statistical pruning to suppress noise, and iteratively refines FSMs using acceptance thresholds to achieve high coverage and precision. Experiments on large-scale srsRAN debug logs demonstrate scalability to gigabyte–terabyte datasets and show that PHOENIX produces interpretable FSMs that enable visualization and facilitate causal reasoning for complex 5G protocol analysis.¹

CCS Concepts

• Networks → Network measurement.

Keywords

Log Analysis, Finite State Machine Inference, 5G Protocols

ACM Reference Format:

Qingyuan Zheng¹, Ushasi Ghosh¹, Ali Mamaghani¹, Srinivas Shakkottai², and Dinesh Bharadia¹. 2025. PHOENIX: Protocol State Machine Recovery from Events in 5G Logs. In *Proceedings of The 16th ACM Wireless of the Students, by the Students, and for the Students Workshop (S3 '25)*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3737898.3769040>

1 Introduction

Modern network systems generate massive volumes of log data and signaling messages during operation. These logs not only contain sufficient detail to infer protocol implementation specifics, but are also indispensable for key functions such as health monitoring and the time-consuming task of root cause failure analysis. However, in multi-vendor network architectures, understanding and evaluating these logs and message flows is highly challenging, as source code is typically unavailable and documentation rarely provides adequate protocol details. Consequently, analyzing logs has

¹Our code is in github link: <https://github.com/ucsdwscng/PHOENIX.git>



This work is licensed under Creative Commons Attribution International 4.0. S3 '25, November 4-8, 2025, Hong Kong, China
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1975-2/2025/11.
<https://doi.org/10.1145/3737898.3769040>

become one of the most practical means of understanding protocol realization. However, understanding the operational dynamics of 5G systems requires more than simply observing raw logs; it demands uncovering the hidden, protocol-driven state transitions that govern network behavior.

As illustrated in ??, the essence of a communication system is defined by the state machines specified in its protocols. To reveal the concrete implementation of these protocols, this work proposes an automated pipeline that infers finite state machines (FSMs) from system operation logs, thereby reconstructing the underlying protocol logic defined in 3GPP standards. By making these latent state transitions explicit, the approach not only enhances the interpretability of system behavior but also enables downstream tasks such as performance diagnosis, fault localization, and adaptive optimization.




Figure 1: From 3GPP Protocols to Hidden State Machines: A Log-Driven Inference Paradigm.

Research on inferring FSMs from execution logs has been active for over a decade. MINT [6] effectively extracts system-level FSMs from small-scale logs, particularly in distributed systems such as HDFS. SCALER [4] extends MINT to component-level inference and heuristically merges these models into system-level FSMs, improving scalability and handling datasets of around 30,000 log entries. However, these methods struggle with the gigabyte- or terabyte-scale logs of communication systems. AutoModel [1] synthesizes FSMs from SoC communication traces by extracting event triplets (source, destination, message) from structured protocol interactions, segmenting them into independent threads via attributes such as module roles and transaction IDs, and enabling precise modular FSM construction. Validated on large-scale SoC datasets, it is conceptually applicable to communication networks with similar concurrency patterns. We adapt and extend this method for FSM modeling from 5G log data.

We propose **PHOENIX**, a fully automated pipeline capable of deriving component-level finite state machines directly from raw network logs. PHOENIX can be applied to any logging system that meets the following two conditions: (1) each log entry contains an explicit network component tag (e.g., RLC, PDCP, MAC); and (2) each log entry contains a user equipment identifier (e.g., UE ID, RNTI).




Figure 2: PHOENIX Overview.

To evaluate the effectiveness of PHOENIX, we analyzed logs from the widely used 5G over-the-air simulation platform SRSRAN. The results demonstrate that PHOENIX is a non-trivial log visualization tool that can support log-based root cause analysis and downstream LLM reasoning tasks.

Related Work

FSM-based inference methods have been widely adopted for protocol reimplementation. Tu et al. proposed 5GBaseChecker, which employs black-box automata learning to model the signaling messages of 5G baseband control-plane protocol implementations as FSMs[5]. It then identifies input sequences that trigger divergent outputs in the learned FSMs and uses these deviations to efficiently locate relevant security properties defined in the protocol specifications.

Another line of work leverages protocol specification documents directly to infer the state machines they define. Hermes[3], for instance, introduces a neural constituency parser and a protocol-oriented domain-specific language to extract transition components from specification texts, translate them into logical formulas, and compile them into FSMs, thereby structurally and precisely modeling protocol flows and enabling the automated formalization of cellular protocol specifications.

2 PHOENIX Design

The PHOENIX pipeline comprises two main modules: the **Log Event Extraction Module**, and the **FSM Inference Module**. **Log Event Extraction Module** leverages state-of-the-art log parsers to infer log templates and identify the complete set of atomic events. During parsing, PHOENIX first groups log entries by component tag and then sequences events from the same session or device using the UE identifier, ultimately producing a time-ordered series of structured events for state machine inference. **FSM Inference Module** adopts AutoModel’s causal inference paradigm in an unsupervised, passive manner to generate per-component finite state models.

Log Event Extraction for FSM. In log-based causal inference, a “log event” is a structured abstraction of raw messages, consisting of a template that captures the fixed textual pattern (e.g., “UE attach request received”), a timestamp for temporal ordering, and parameter values extracted from variable segments (e.g., UE ID, sequence number, error code). Representing logs as *type + time + attributes* provides a consistent and analyzable input for causal inference.

To obtain such structured events, log parsing converts unstructured messages into templates by tokenizing text, normalizing variable elements (e.g., IP addresses, IDs), separating constant and variable components, and clustering similar structures based on token

sequence or pattern similarity for large-scale aggregation. In this work, we adopt **Drain** [2] as the primary parsing method. Drain organizes templates into a fixed-depth parse tree: messages are routed by length, filtered by leading tokens, and matched at leaf nodes—updating matches with wildcards or creating new templates when unmatched. This one-pass streaming process enables high-throughput log parsing without offline training.

Once logs are structured, they can be naturally modeled as sequences of discrete events, which aligns with the formalism of an FSM. An FSM models the behavior of a discrete event system and, in its deterministic form, is defined as a tuple: $M = (Q, \Sigma, \delta, q_0)$, where Q is the finite set of states, Σ is the finite event alphabet, $\delta : Q \times \Sigma \rightarrow Q$ is the transition function, and $q_0 \in Q$ is the initial state. Given an event sequence $\sigma = e_1 e_2 \dots e_n \in \Sigma^n$, the FSM generates a unique state path: $q_0 \xrightarrow{e_1} q_1 \xrightarrow{e_2} \dots \xrightarrow{e_n} q_n$, with $q_i = \delta(q_{i-1}, e_i)$.

In log analysis, each parsed log line corresponds to an event $e_i \in \Sigma$ with an associated timestamp, forming a trace σ that reflects the hidden execution paths of the underlying FSM. FSM inference from logs thus aims to recover the observable event set Σ , the hidden state set Q , the initial state q_0 , and the transition relation δ .

FSM Inference. The FSM Inference Module reconstructs protocol behavior from raw log data \mathcal{L} through a sequence of structured steps, fully parameterized by the confidence threshold θ , the acceptance threshold τ , and the maximum number of iterations T .

Directed Acyclic Graph (DAG) Construction: We first parse the raw log stream \mathcal{L} into an ordered sequence of events, assigning each unique event a distinct identifier. For each event node s , we compute its **NodeSupport** $N(s)$, defined as the total number of occurrences of s in \mathcal{L} . For each ordered pair of adjacent events (a, b) , we calculate the **EdgeSupport** $E(a, b)$, representing the number of times a is immediately followed by b . Two directional confidence measures are defined as $FC(a, b) = E(a, b)/N(a)$, $BC(a, b) = E(a, b)/N(b)$, where $FC(a, b)$ denotes the probability of b following a , and $BC(a, b)$ denotes the probability of a preceding b . Using the set of nodes V and all observed edges, we construct a directed acyclic graph (DAG) $G = (V, E)$.

Graph Pruning: To eliminate noisy or weakly supported transitions, we compute for each edge $(a, b) \in E$ the average confidence score: $\text{Conf}_{\text{avg}}(a, b)$. If $\text{Conf}_{\text{avg}}(a, b) < \theta$, the edge is removed from G . This process retains only transitions with strong bidirectional statistical support, improving clarity and reducing complexity for subsequent FSM extraction.

Initial FSM Extraction: From each connected component of the pruned DAG, we extract the longest node-disjoint path P_i to maximize coverage of distinct nodes. These paths are concatenated, in

