ProvGuard: Detecting SDN Control Policy Manipulation via Contextual Semantics of Provenance Graphs

Ziwen Liu, Jian Mao, Jun Zeng, Jiawei Li, Qixiao Lin, Jiahao Liu,
Jianwei Zhuge, Zhenkai Liang

Network and Distributed System Security (NDSS) Symposium 2025

NUS
National University
of Singapore

Software-Defined Networking

- Software-Defined Networking (SDN) separates network control from forwarding devices into the control plane

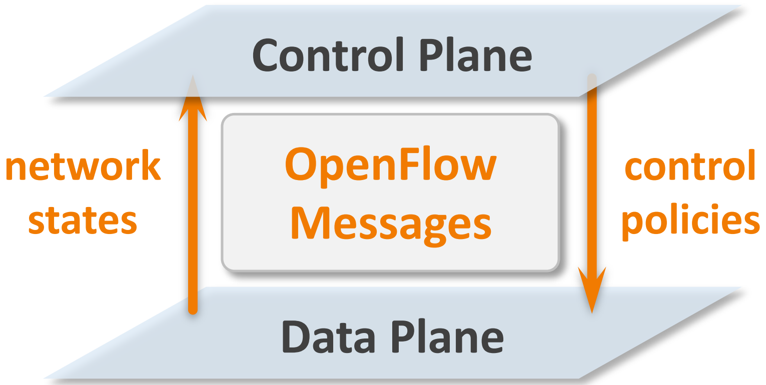

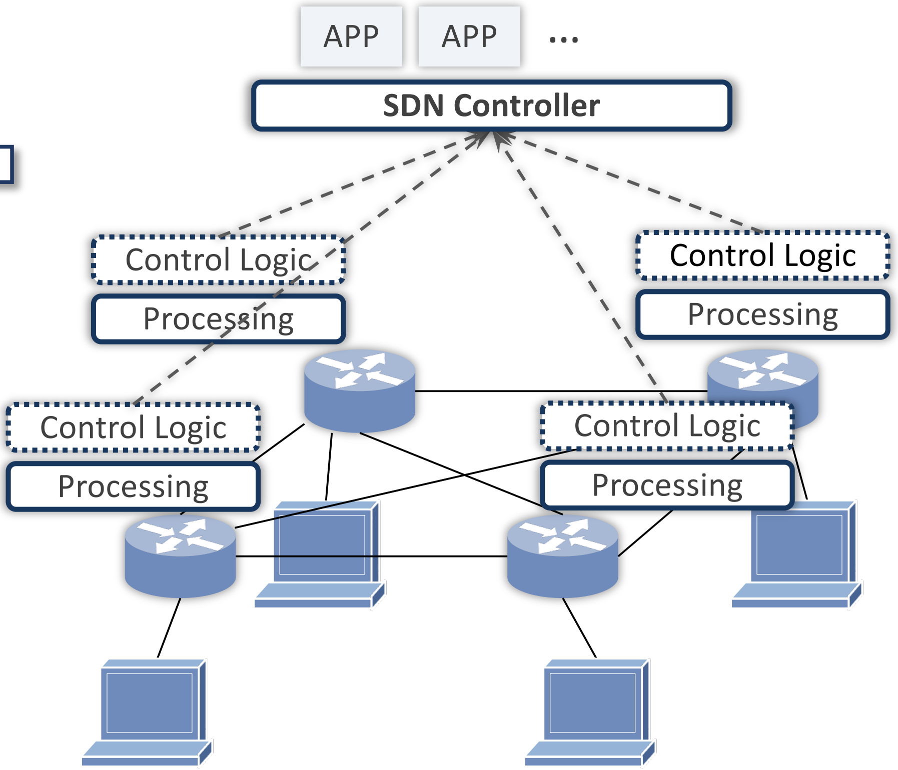

Software-Defined Networking

- Software-Defined Networking (SDN) separates network control from forwarding devices into the control plane

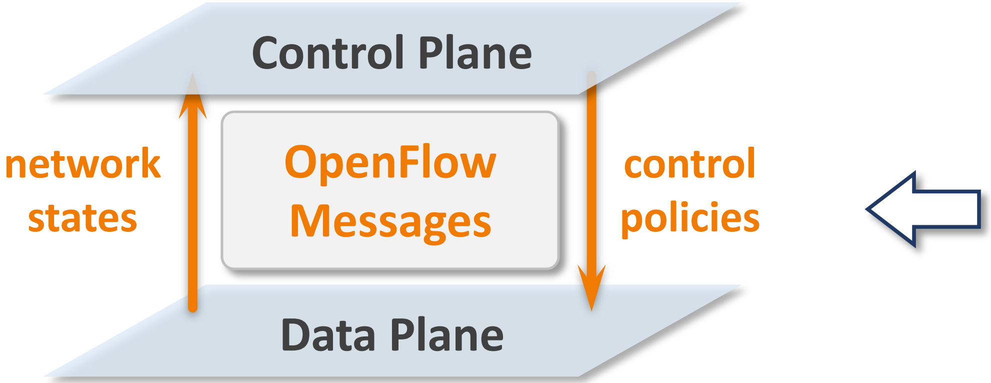

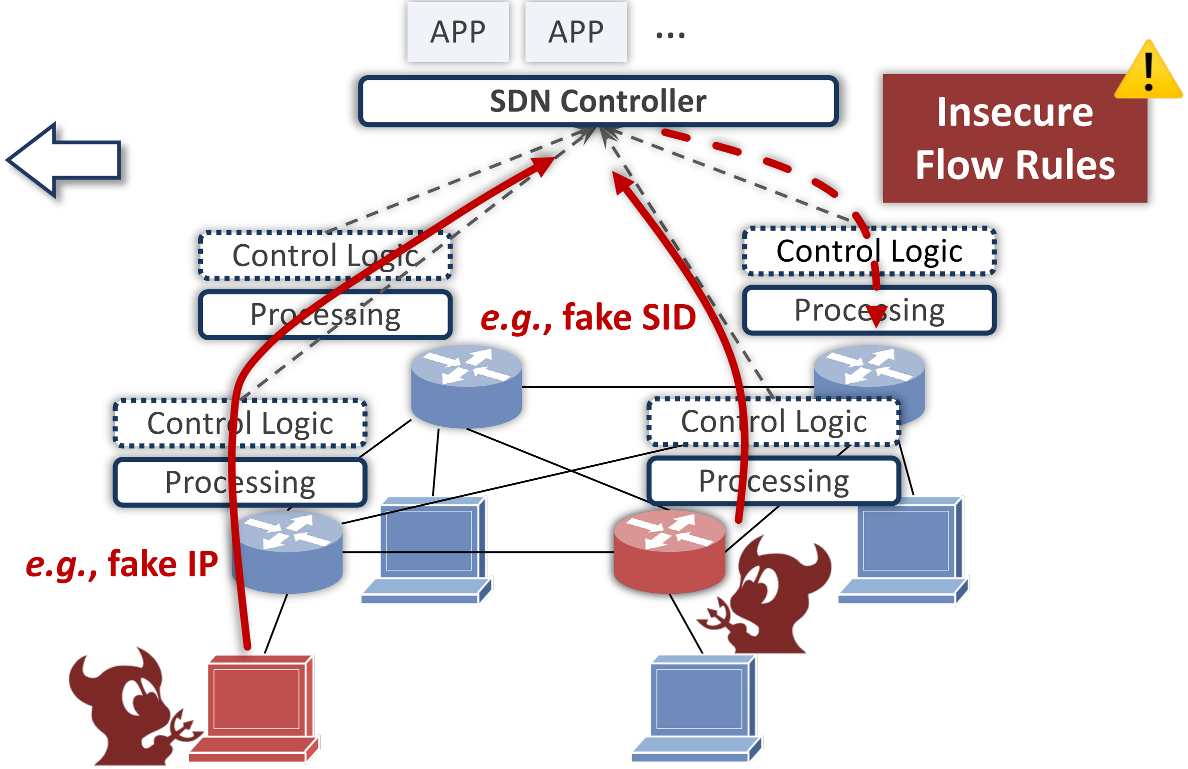

- Control Plane Manipulation
  - modifies or deactivates network forwarding and security policies

Related Work on CPM Defense

Anomaly Detection

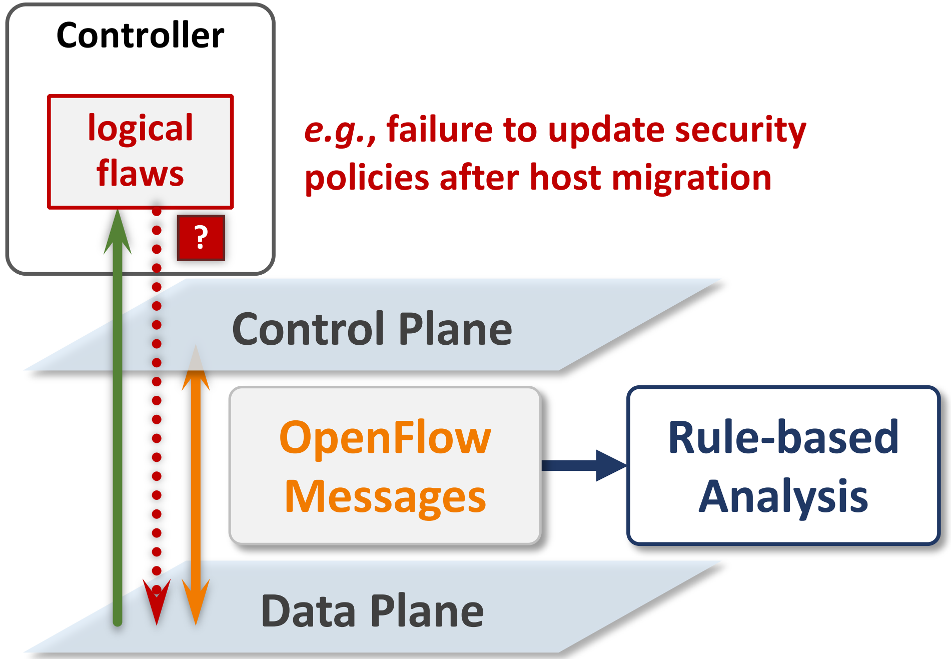

e.g., failure to update security policies after host migration

- Fail to detect attacks exploiting logic flaws by *normal* data-plane operations

Related Work on CPM Defense

Anomaly Detection

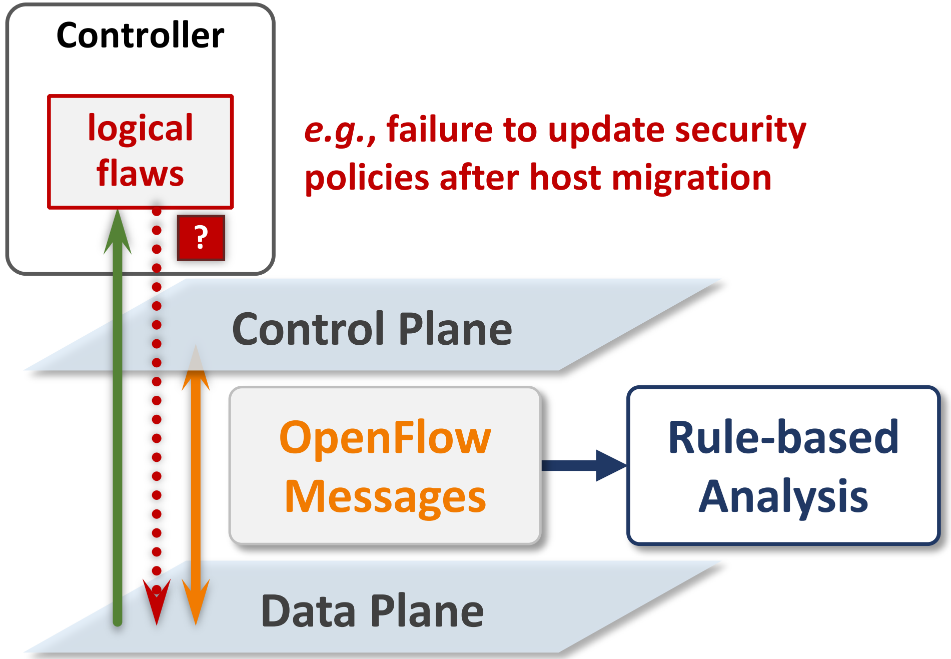

- Fail to detect attacks exploiting logic flaws by *normal* data-plane operations

Policy Verification

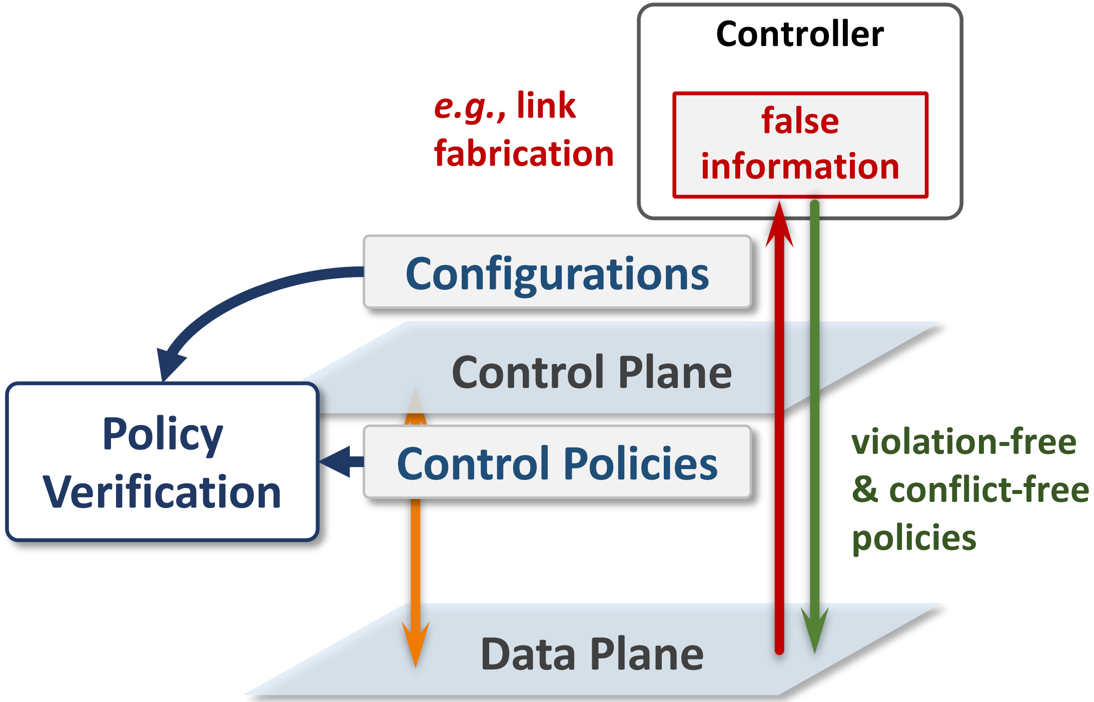

- Fail to prevent CPM attacks that *do not* cause policy violations or conflicts

Motivation

- *Controller operations provide direct insights into network state changes and their impact on control decisions*

Motivation

- *Controller operations provide direct insights into network state changes and their impact on control decisions*

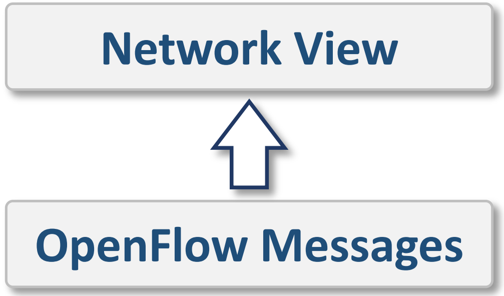

Network View

OpenFlow Messages

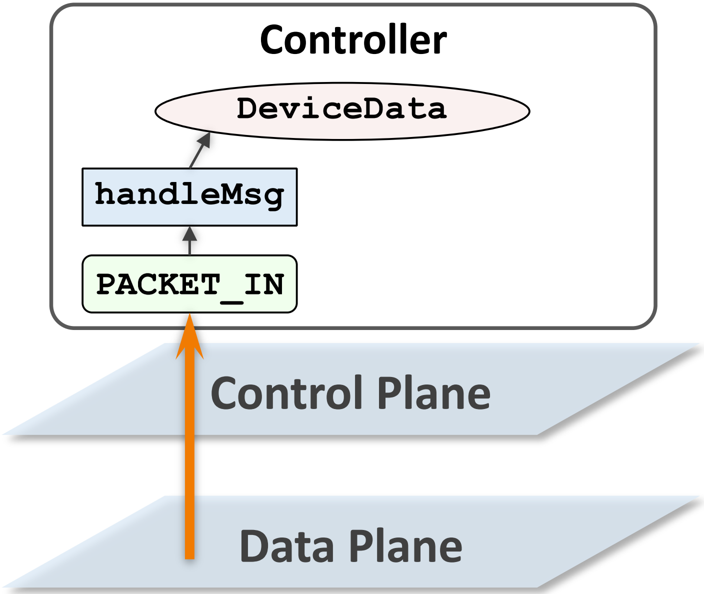

Controller

DeviceData

handleMsg

PACKET_IN

Control Plane

Data Plane

Motivation

- *Controller operations provide direct insights into network state changes and their impact on control decisions*

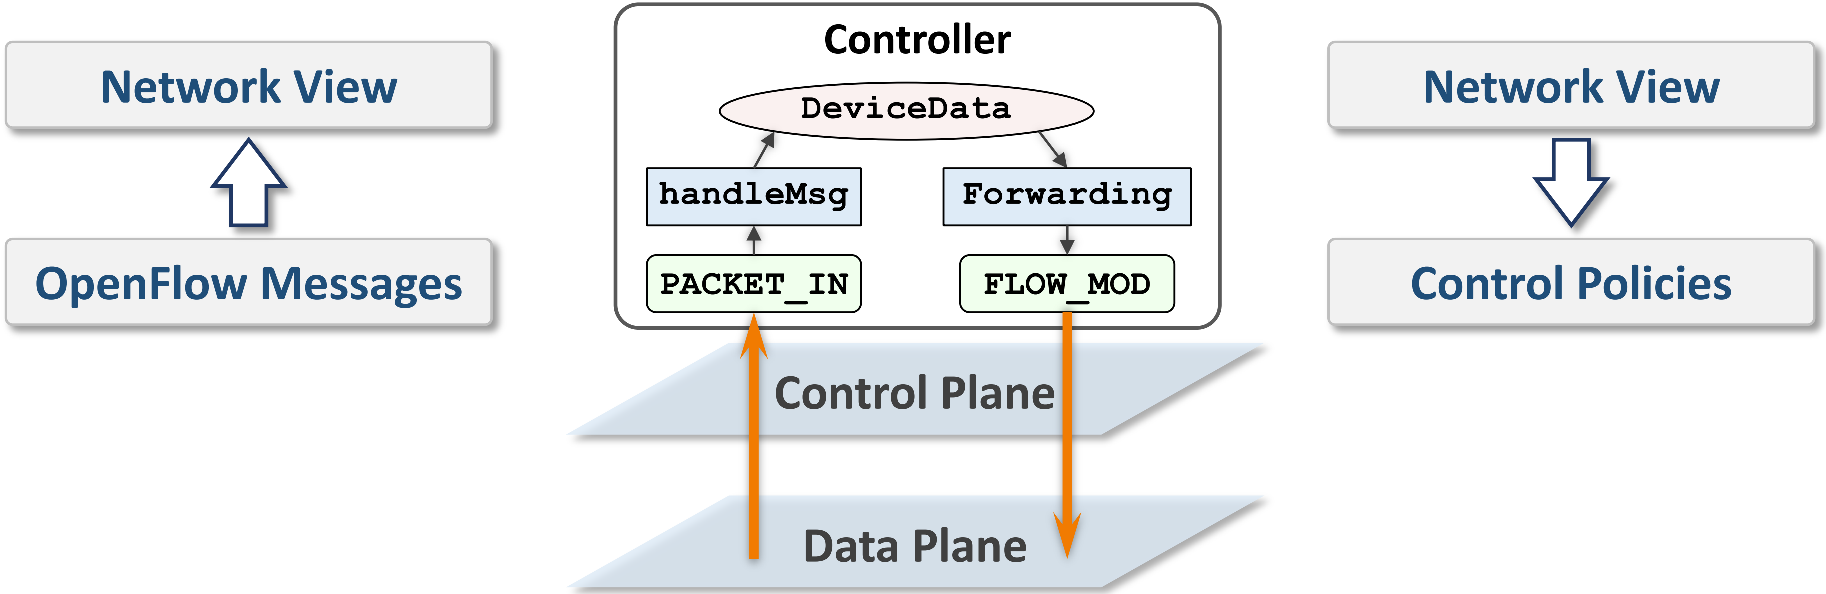

Network View

OpenFlow Messages

Controller

DeviceData

handleMsg

Forwarding

PACKET_IN

FLOW_MOD

Control Plane

Data Plane

Network View

Control Policies

Motivation

- *Controller operations provide direct insights into network state changes and their impact on control decisions*

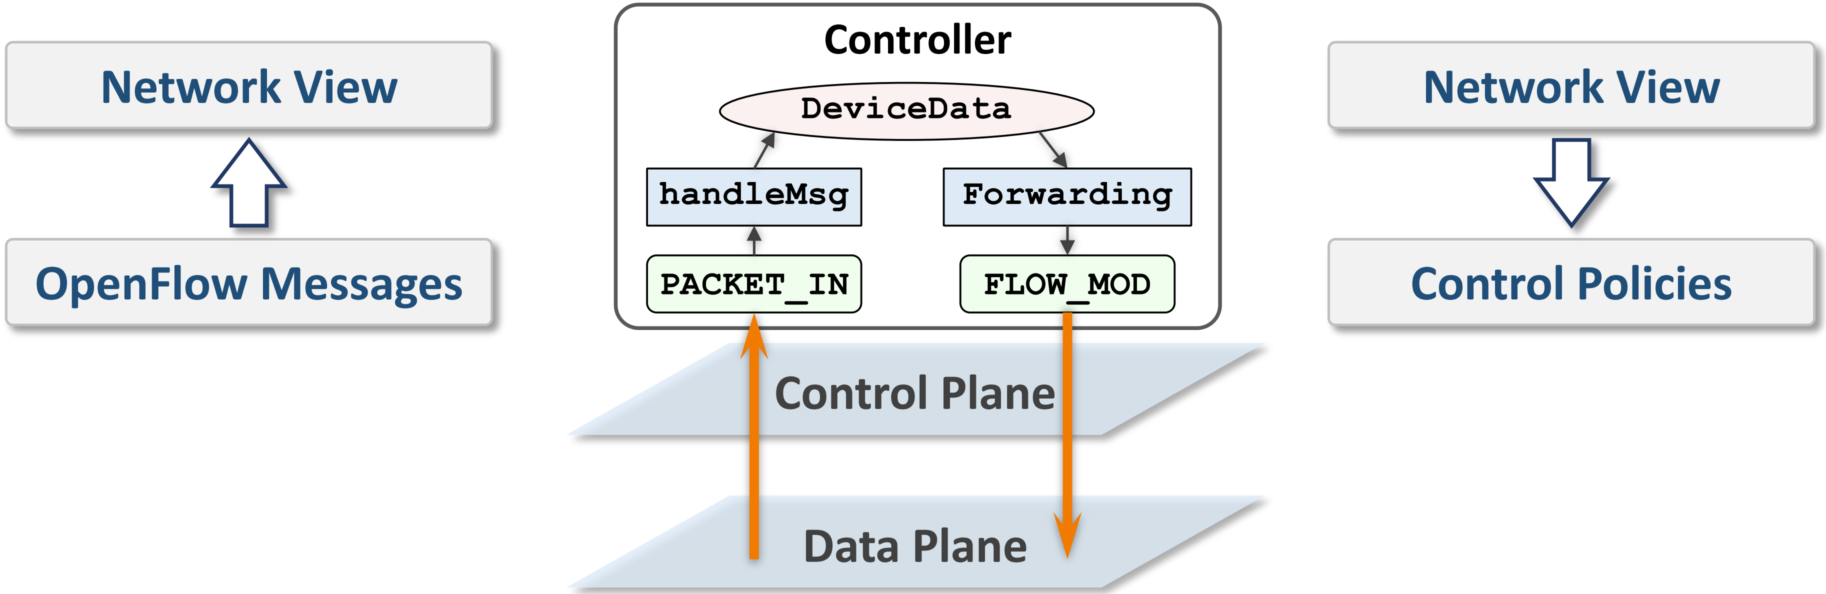

- Use ***provenance graph*** to describe the ***causal dependencies*** between entities (function, data, thread, event…) in control plane

Main Challenge

- *How to extract features of long-term CPM attacks from control plane provenance graph*

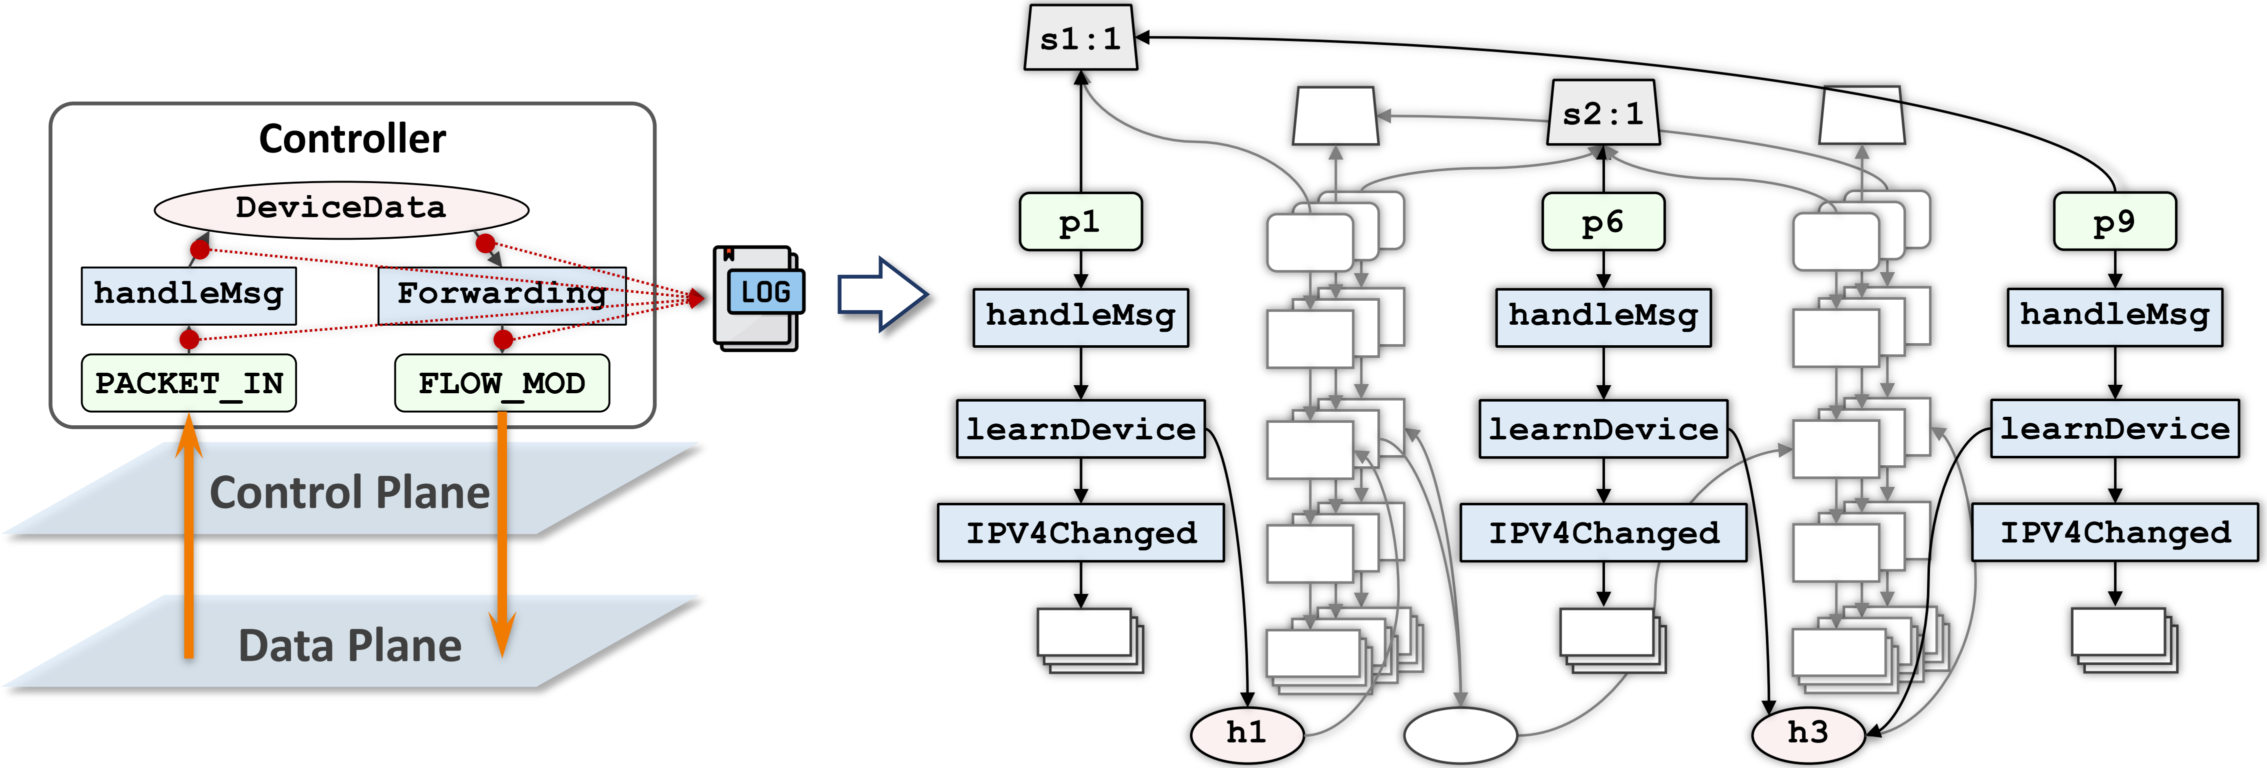

Main Challenge

- *How to extract features of long-term CPM attacks from control plane provenance graph*

Unclear behavior boundaries

Behavior subgraph partition

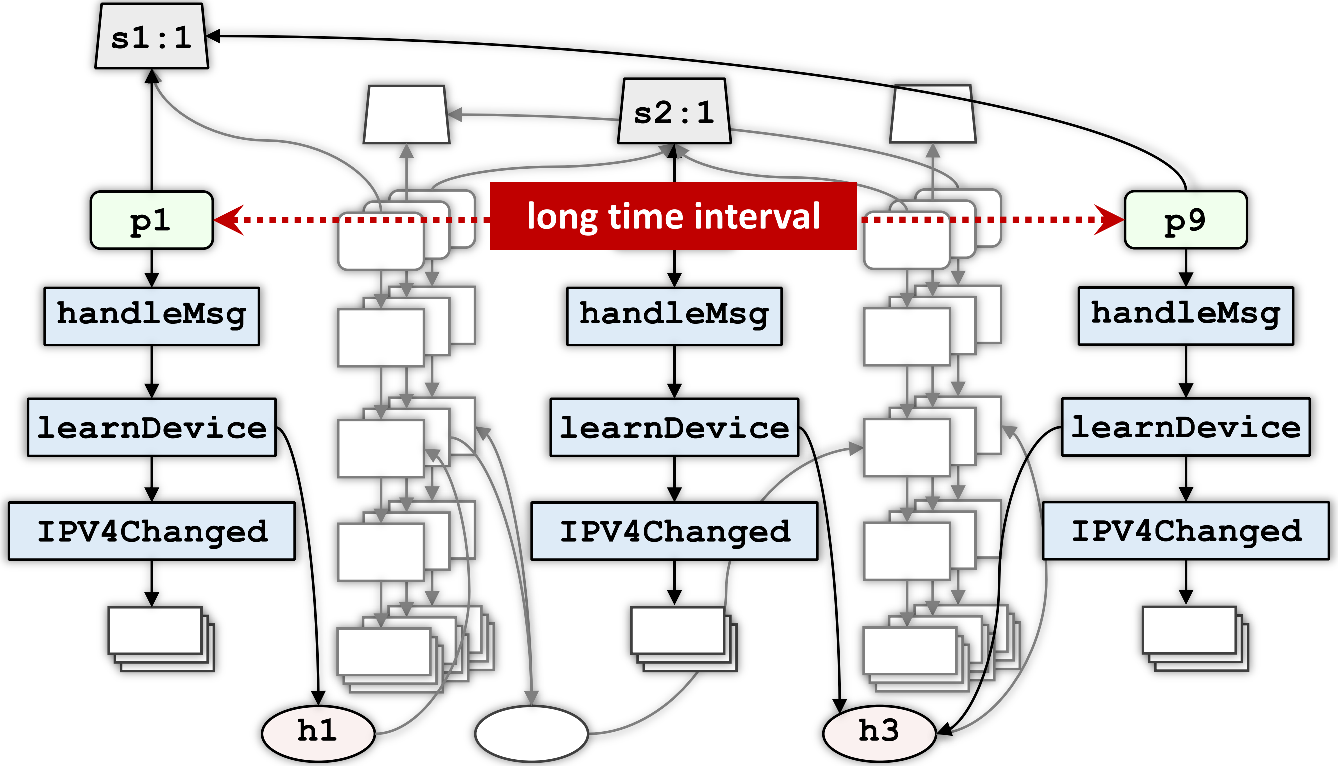

Main Challenge

- *How to extract features of long-term CPM attacks from control plane provenance graph*

Unclear behavior boundaries

✗ Behavior subgraph partition

✓ Multi-stage feature association

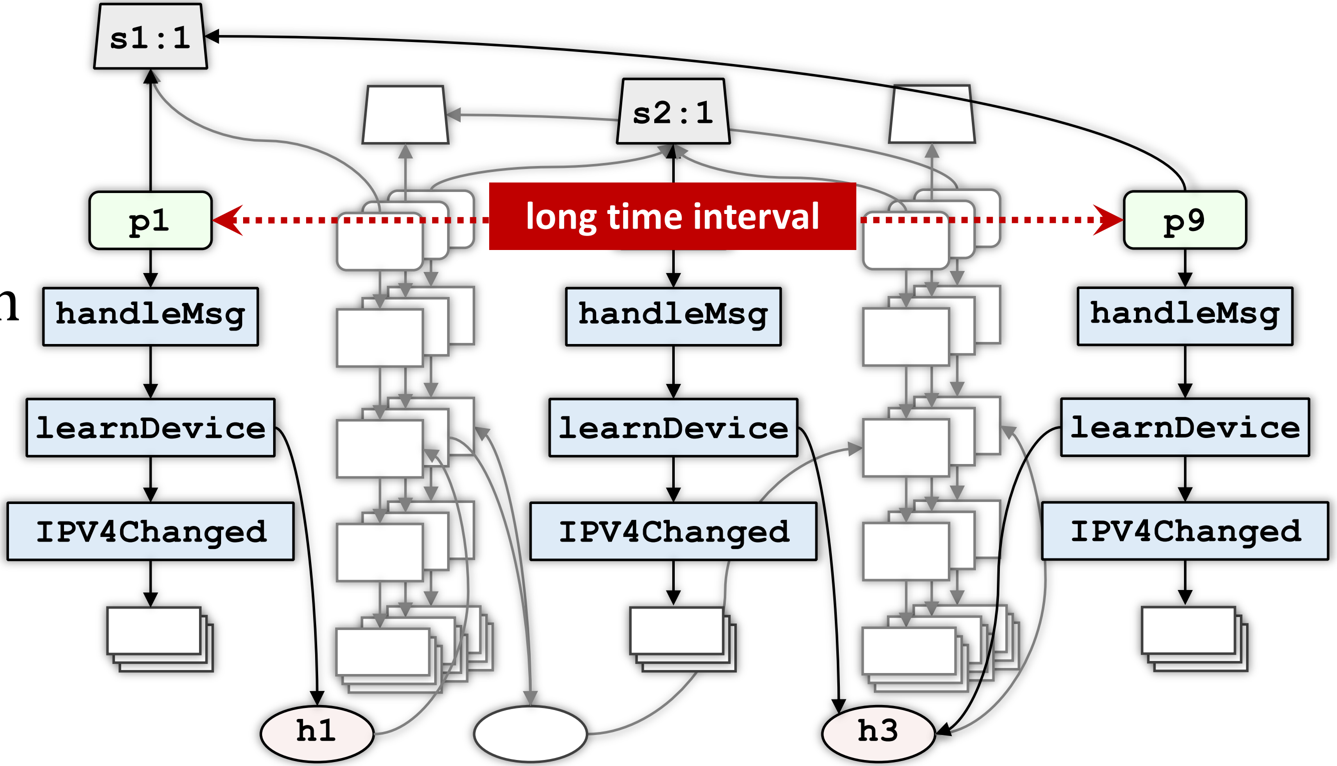

Main Challenge

- *How to extract features of long-term CPM attacks from control plane provenance graph*

Unclear behavior boundaries

✗ Behavior subgraph partition

✓ Multi-stage feature association

Neighborhood similarity

Fixed-hop graph features

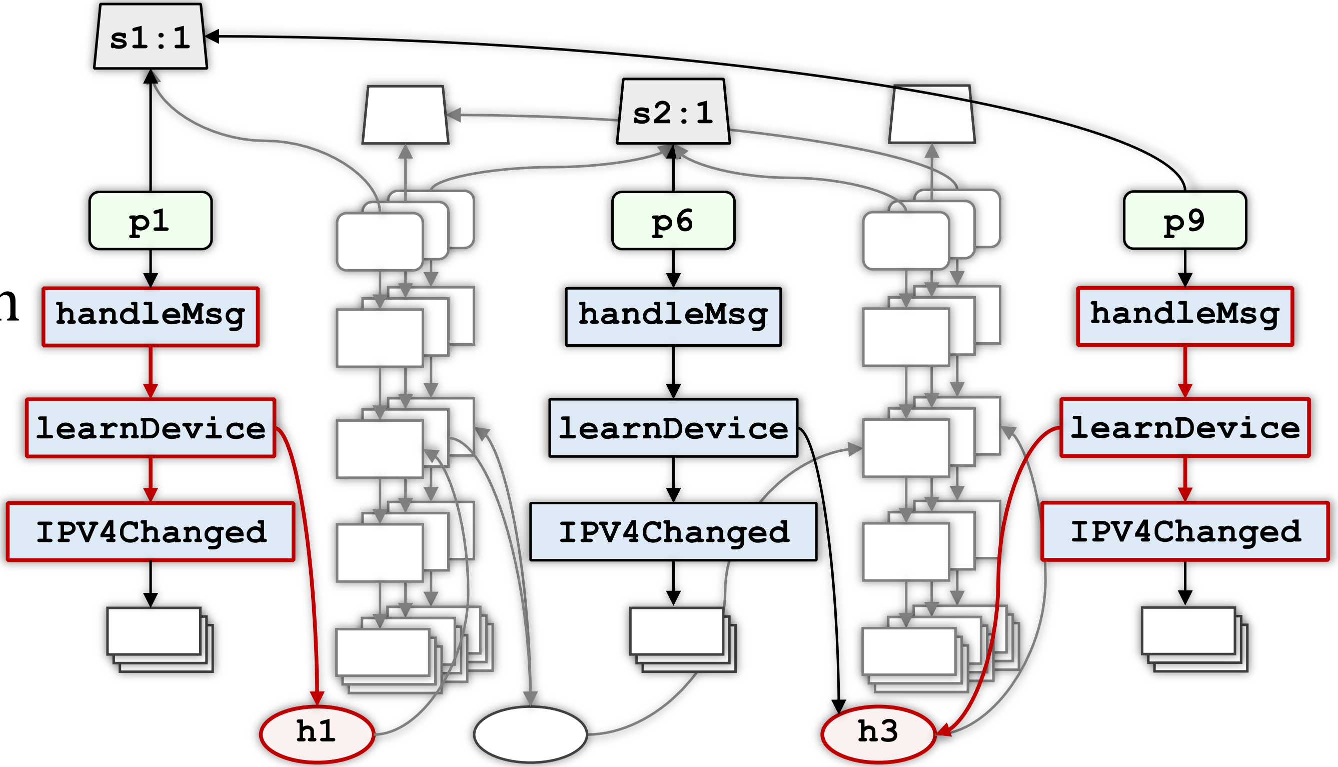

Main Challenge

- *How to extract features of long-term CPM attacks from control plane provenance graph*

Unclear behavior boundaries

✗ Behavior subgraph partition

✓ Multi-stage feature association

Neighborhood similarity

✗ Fixed-hop graph features

✓ Contextual semantics in ***paths***

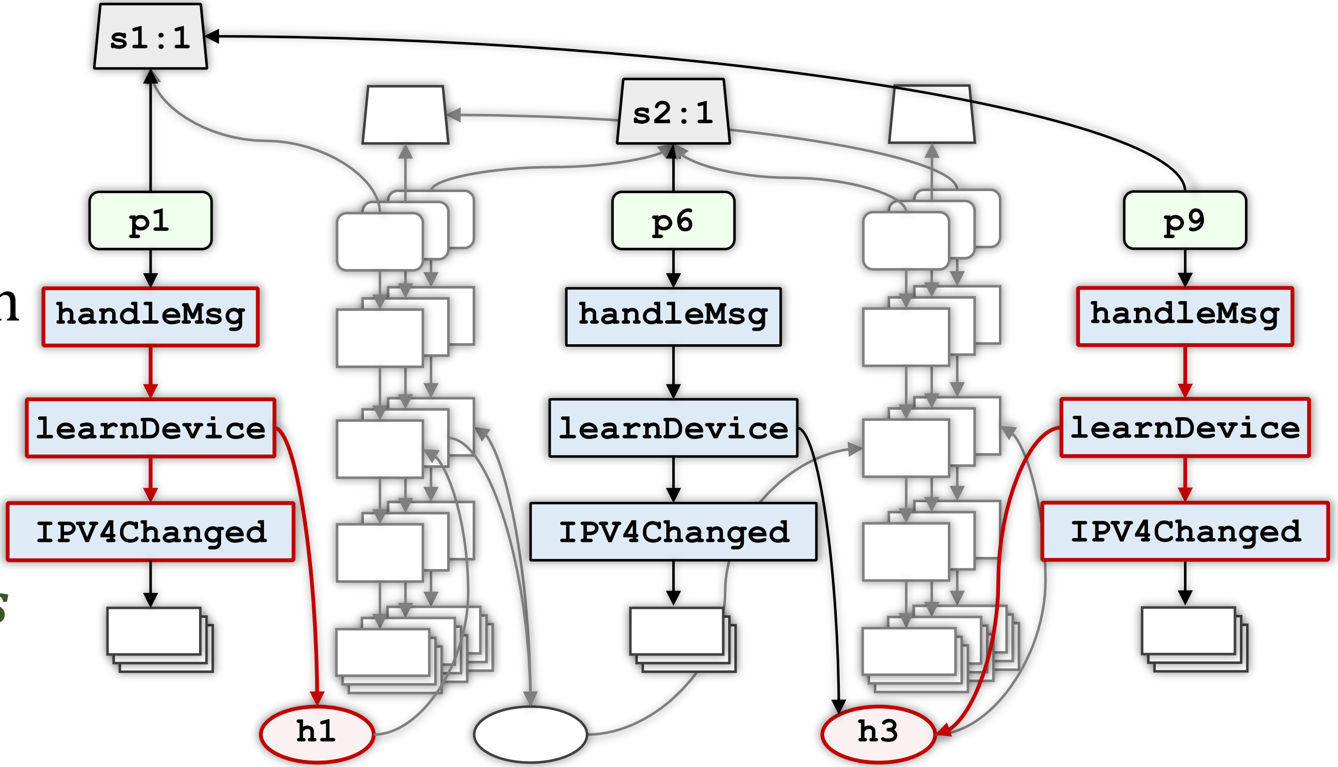

ProvGuard: Detecting CPM on Provenance Graph

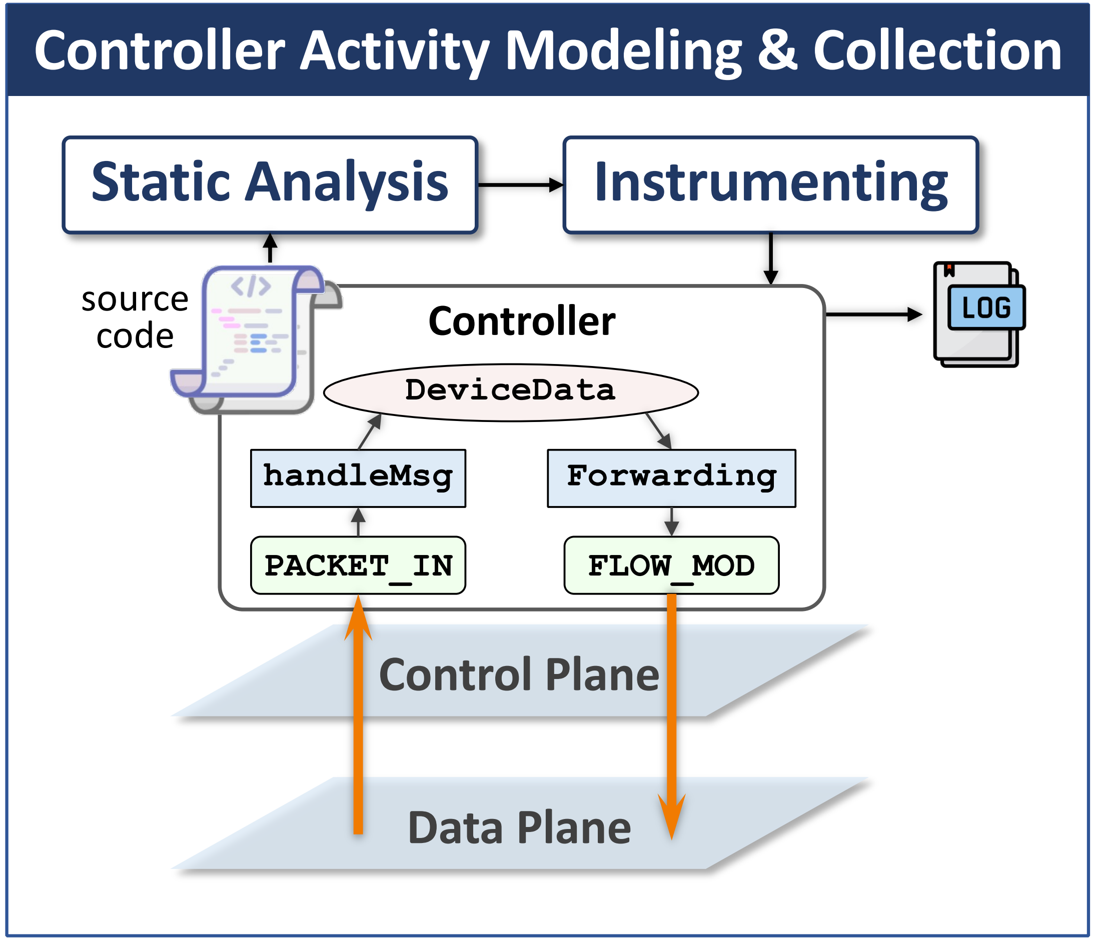

ProvGuard: Detecting CPM on Provenance Graph

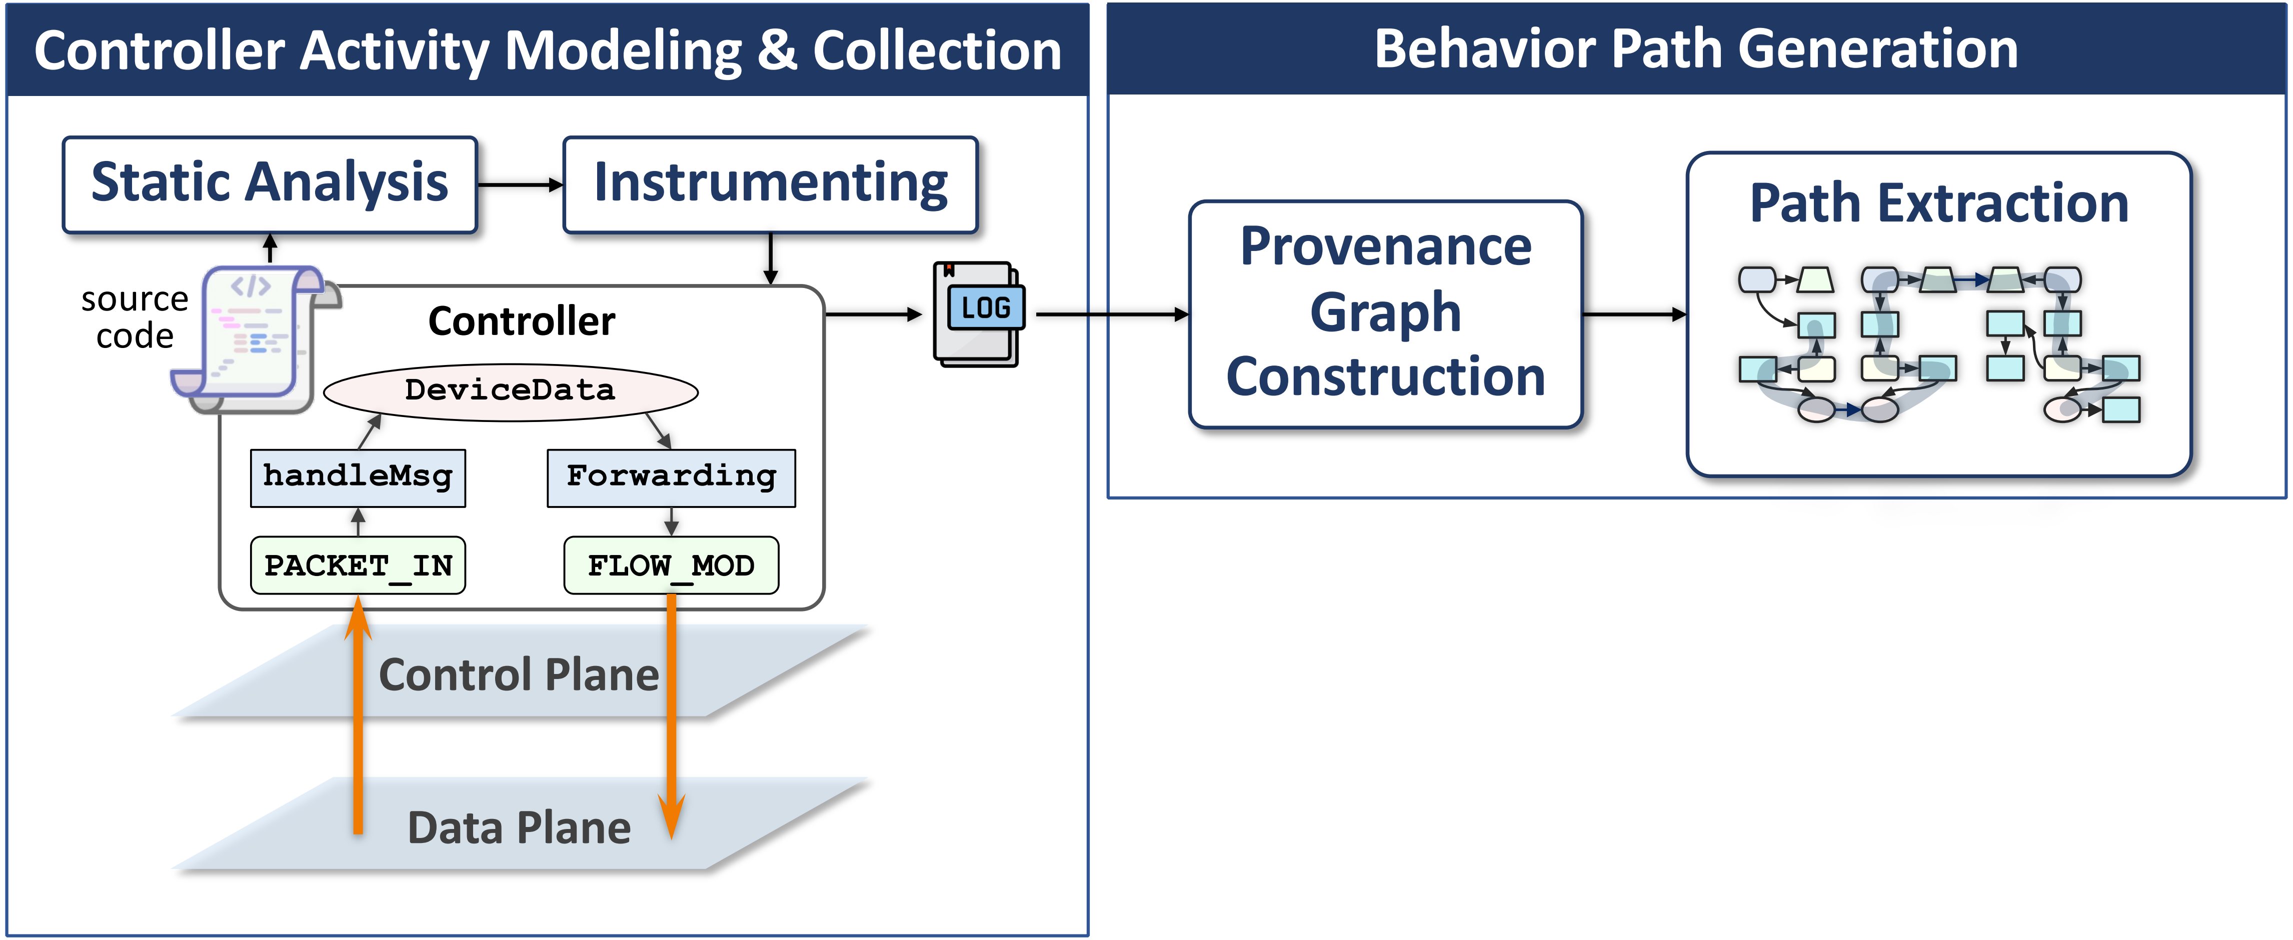

ProvGuard: Detecting CPM on Provenance Graph

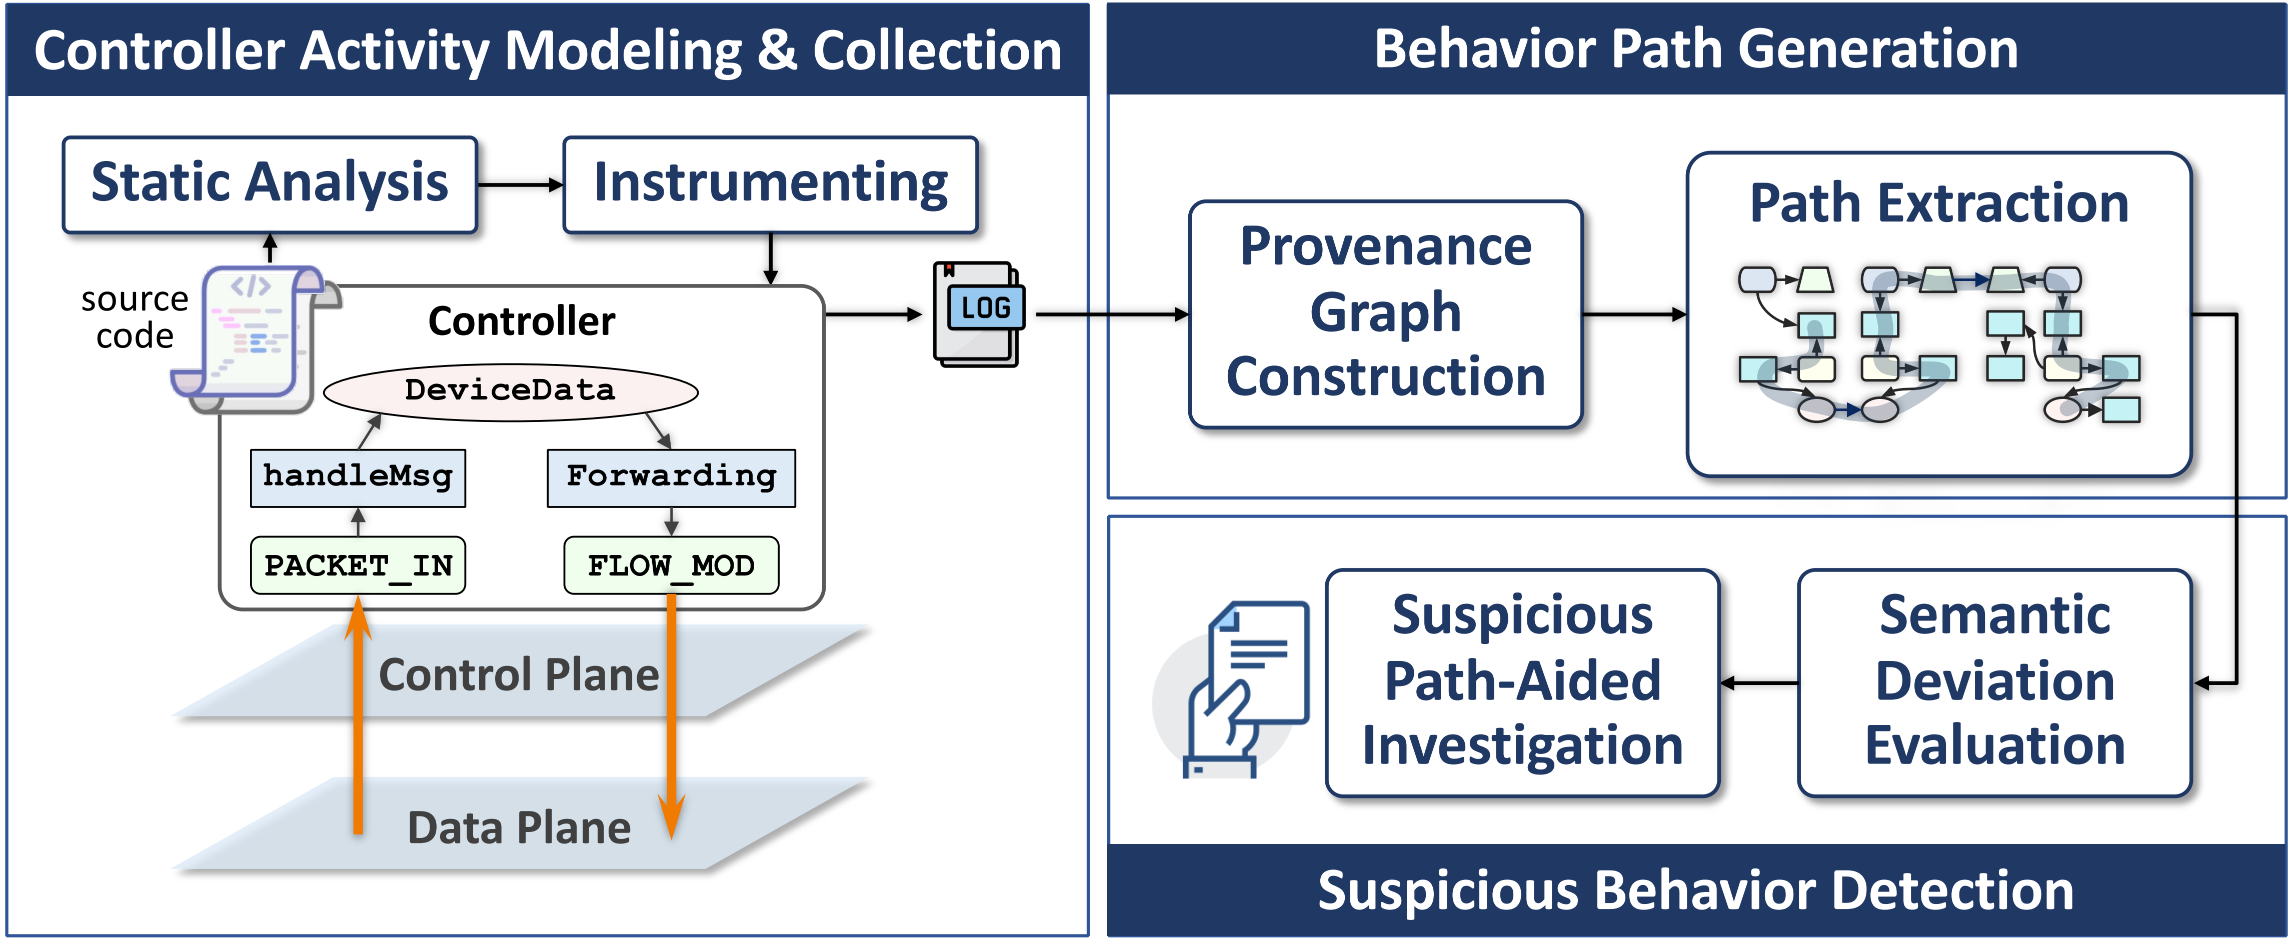

Controller Activity Modeling & Collection

Static Analysis

Instrumenting

source code

Controller

DeviceData

handleMsg

Forwarding

PACKET_IN

FLOW_MOD

Control Plane

Data Plane

LOG

Behavior Path Generation

Provenance Graph Construction

Path Extraction

Semantic Deviation Evaluation

Suspicious Path-Aided Investigation

Suspicious Behavior Detection

Controller Activity Modeling and Collection

- Capture data-plane message's impact on control policies
  - Analyze controller source code from data-plane message handler

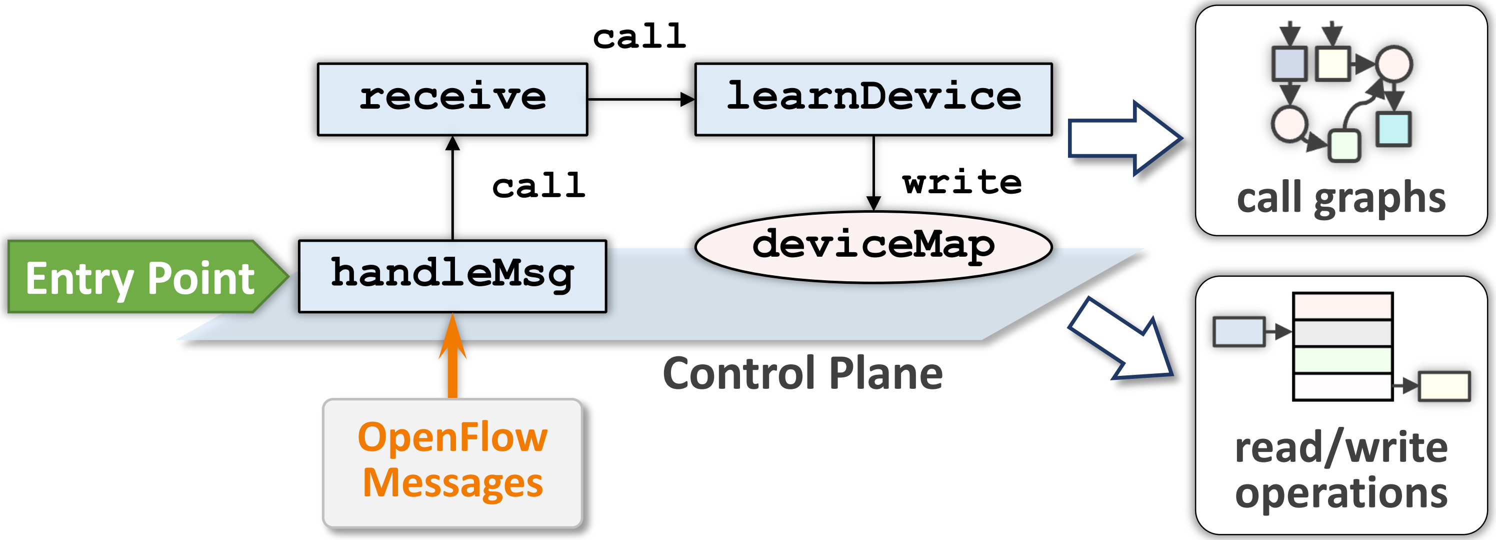

Controller Activity Modeling and Collection

- Capture data-plane message’s impact on control policies
  - Analyze controller source code from data-plane message handler

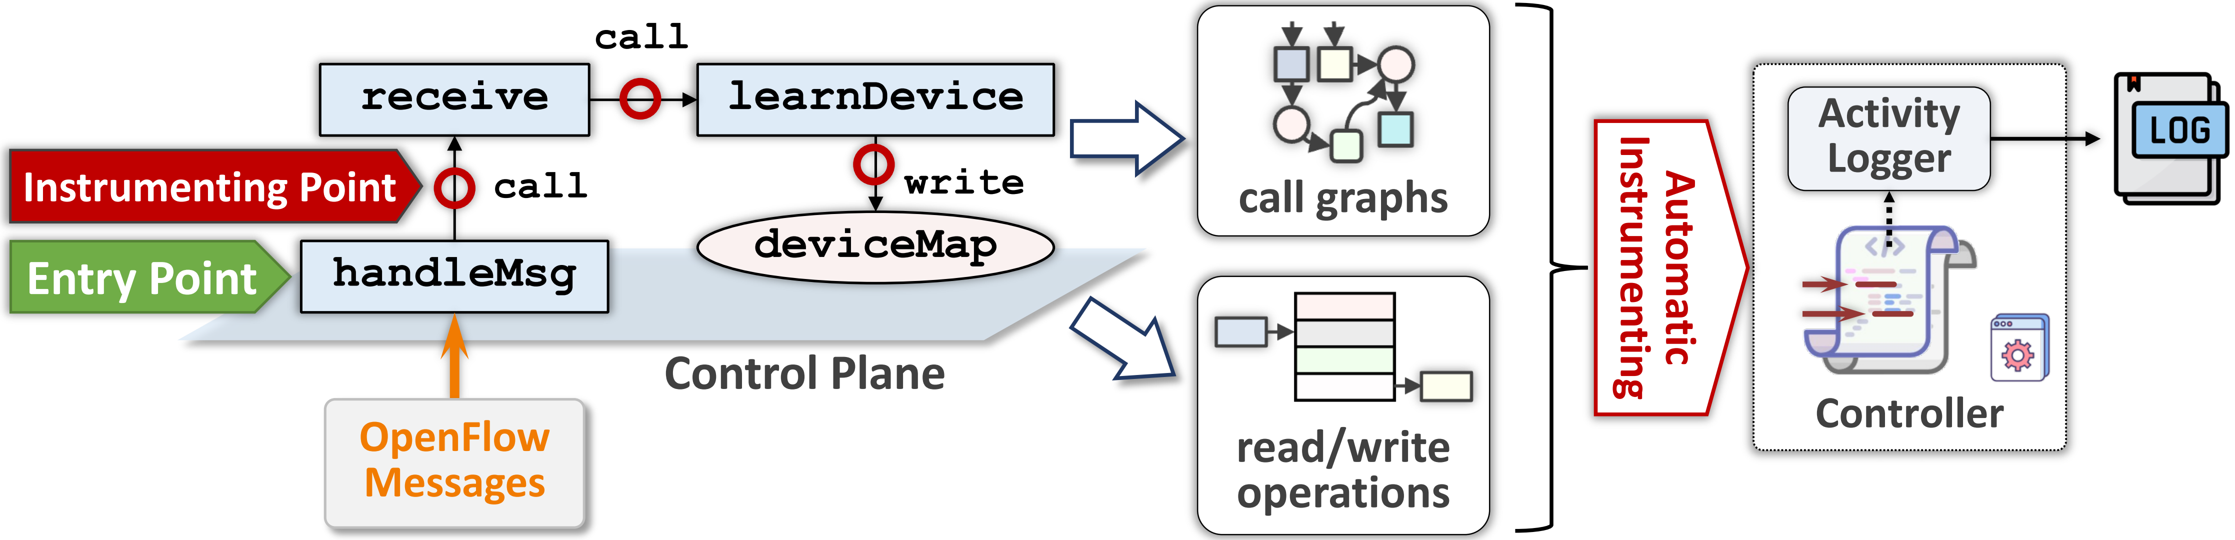

- Log Controller Activities
  - Insert collectors into the controller to record activities

Behavior Path Generation

- Reconstruct and associate execution unit graphs

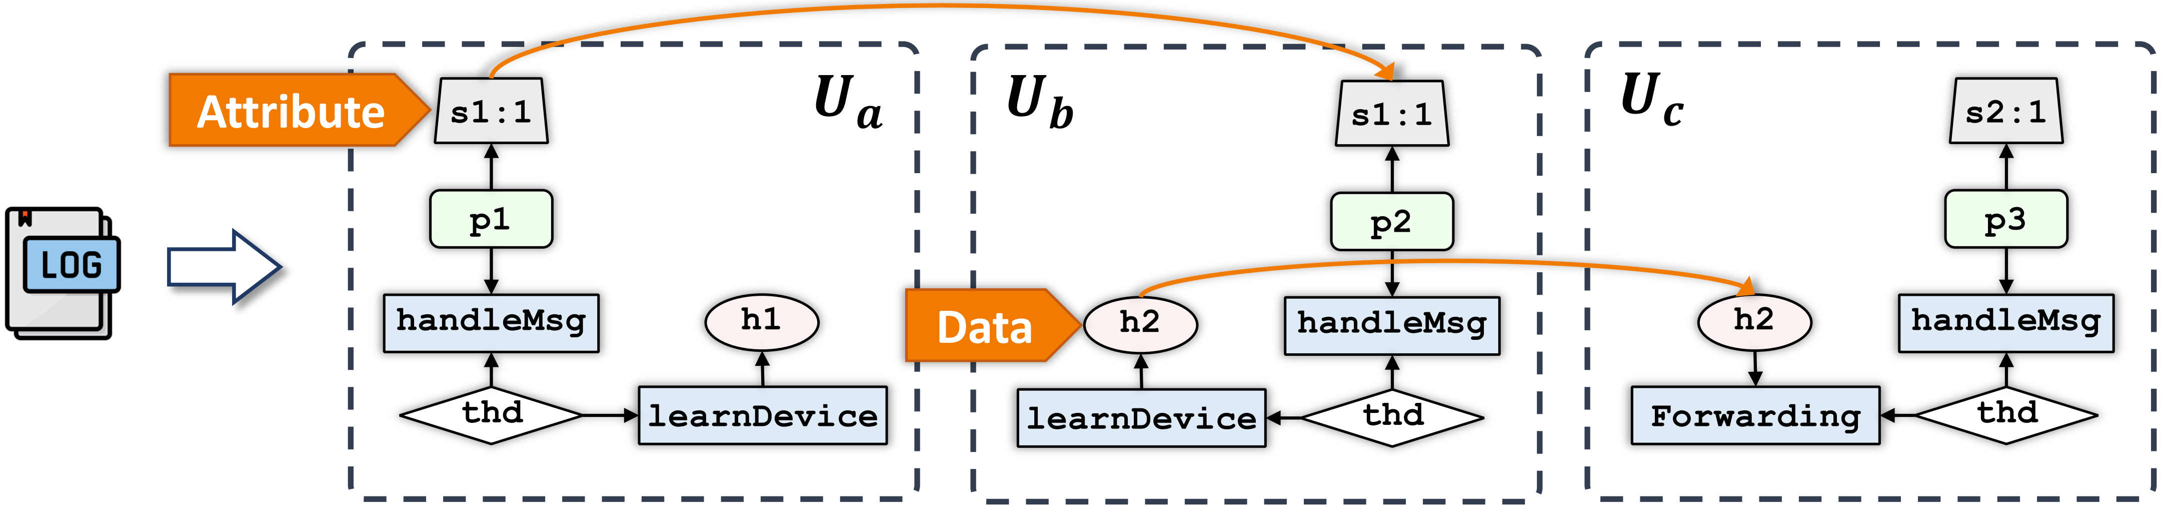

- Reduce redundancy
  - Assess edge/unit importance via inverse document frequency
  - Filter out frequent operations and patterns

Behavior Path Generation

- Reconstruct and associate execution unit graphs

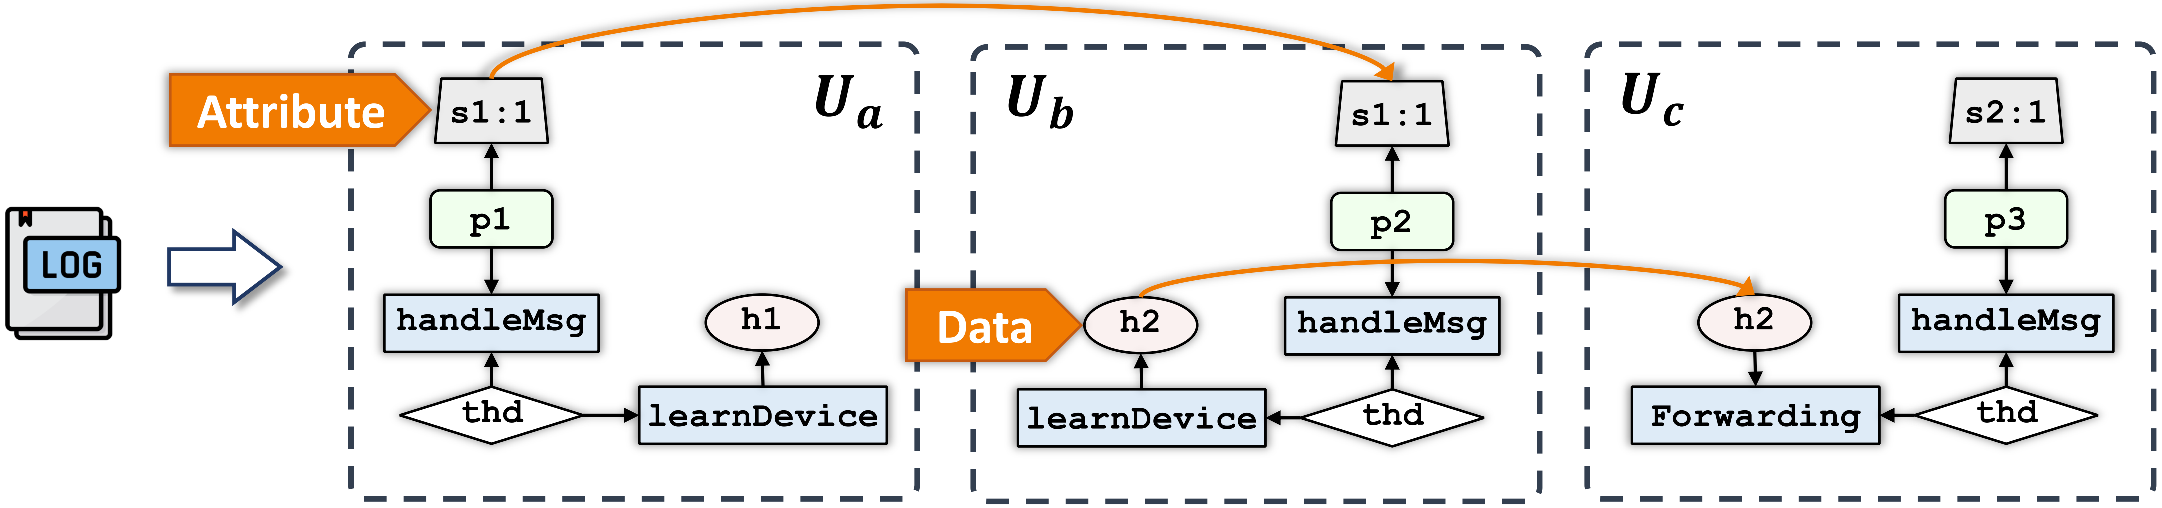

- Extract paths
  - Search sub-paths inside execution unit graphs
  - Associate intra-unit paths via inter-unit edges

Suspicious Behavior Detection

- Long-term CPM detection
  - Multistage feature extraction ⇒ paths spanning execution units
  - Attack-agnostic detection ⇒ contextual semantics learning and deviation identification

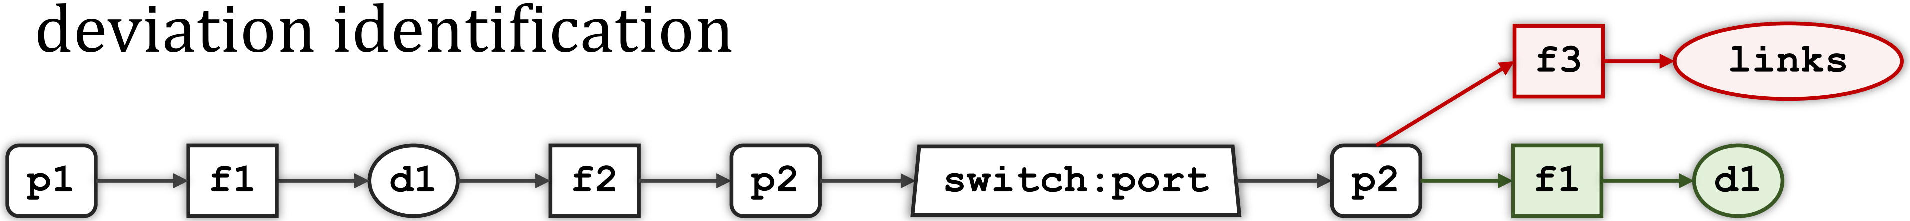

Suspicious Behavior Detection

- Long-term CPM detection
  - Multistage feature extraction ⇒ paths spanning execution units
  - Attack-agnostic detection ⇒ contextual semantics learning and deviation identification

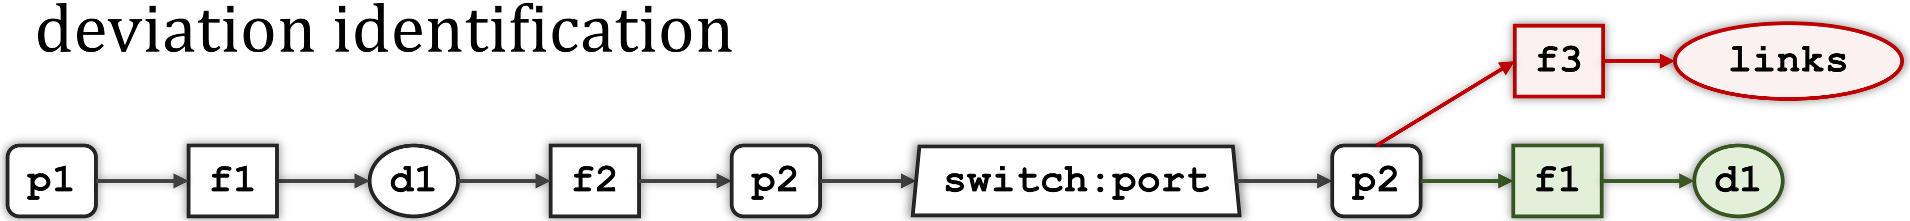

- Semantic Deviation Evaluation
  - Abnormal paths cause larger prediction errors

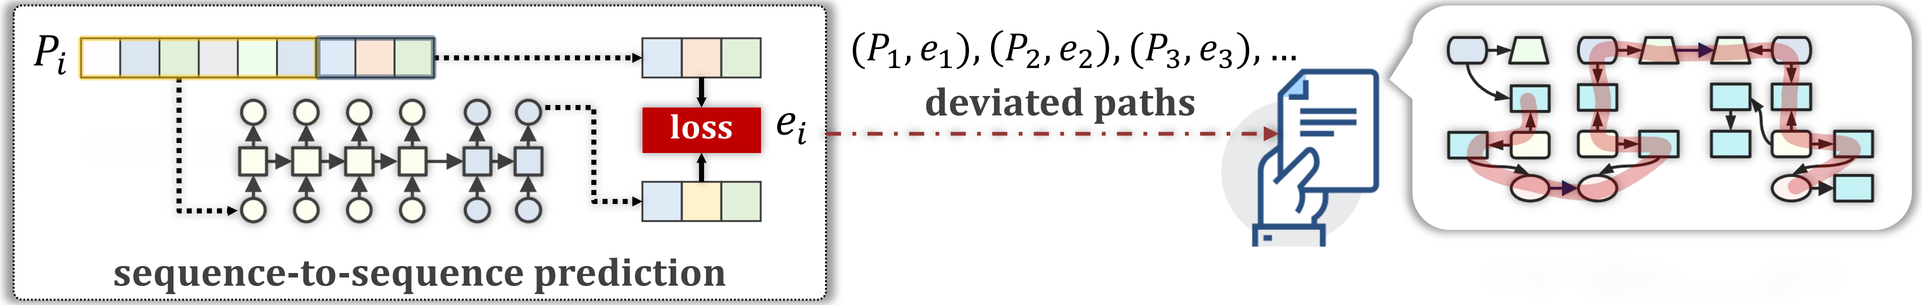
P_i
loss
e_i
sequence-to-sequence prediction
$(P_1, e_1), (P_2, e_2), (P_3, e_3), \ldots$
deviated paths

Evaluation

- Implementation
  - Floodlight (SDN controller), Mininet (network simulation)
  - Data Collection
    - [Normal] Representative host behaviors
    - [Abnormal] Four typical CPM attacks

- Evaluation Aspects
  - How effectively ProvGuard detects CPM attacks?
  - How effective is the redundancy reduction in filtering out noises?
  - How contextual semantics contribute to anomaly detection?
  - How much ProvGuard reduces the manual effort for log auditing?
  - Is the overhead of controller activity collection acceptable?

Effectiveness of CPM detection

- ProvGuard effectively captures long-term CPM features

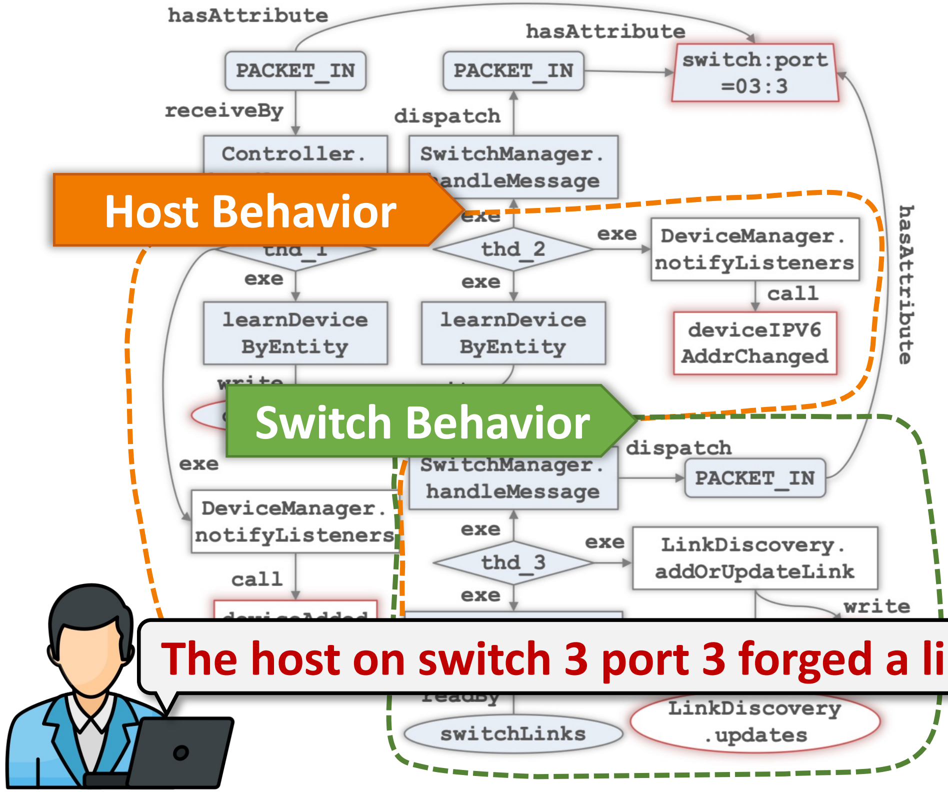

Host Behavior

Switch Behavior

The host on switch 3 port 3 forged a link

Effectiveness of CPM detection

- ProvGuard effectively captures long-term CPM features
- ProvGuard outperforms existing detection approaches in identifying CPM attacks

	Network Identifier Hijacking	Link Fabrication	Access Control Bypass	Switch ID Spoofing
SPHINX	✓	✓		
Veriflow			✓	
PacketChecker	✓			
TopoGuard	✓	✓		
SPV			✓	
FlowChecker			✓	
ProvGuard	✓	✓	✓	✓

* conceptually comparison

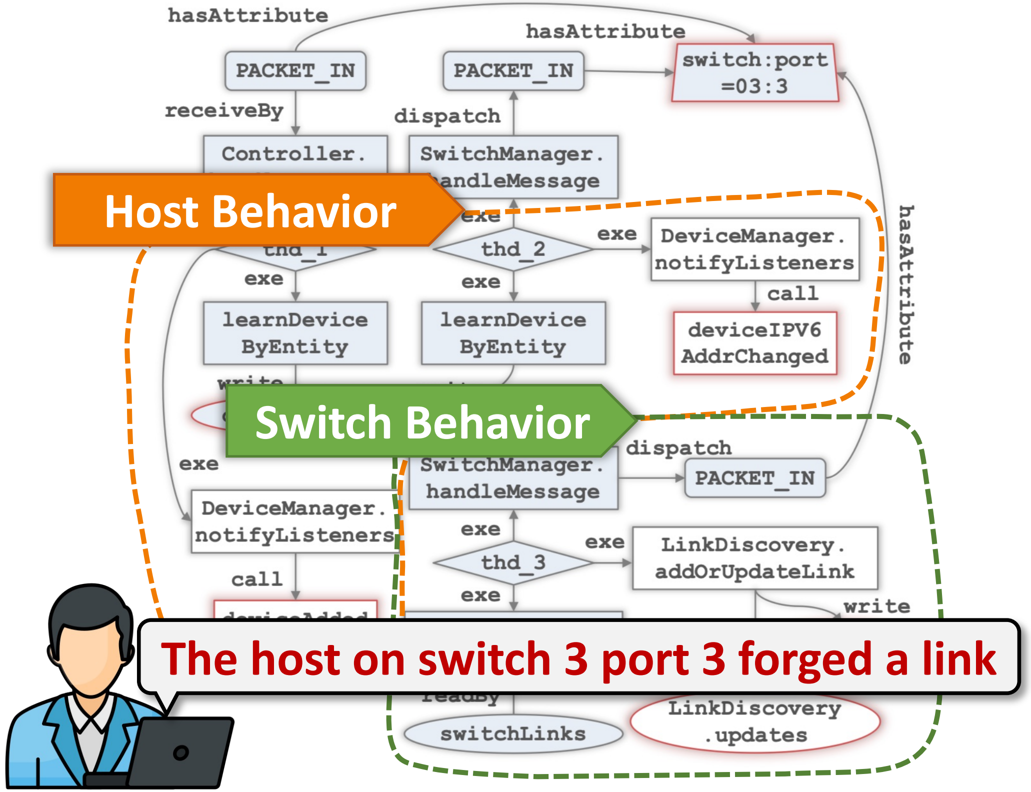

The host on switch 3 port 3 forged a link

Performance & Effect of Context Extraction

- Effectively reduce extracted paths per log entry

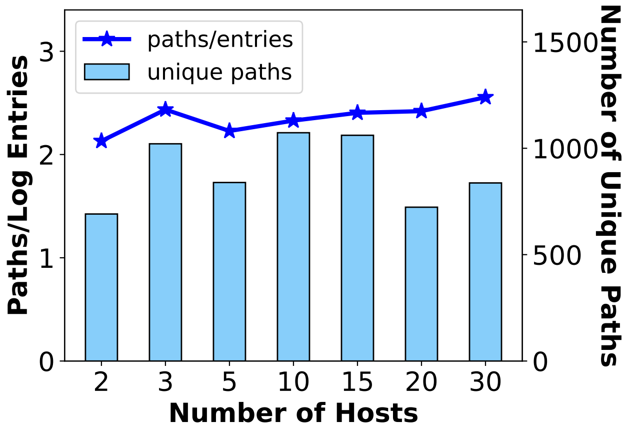

  - The number of unique paths stable regardless of network scales and log volumes

Performance & Effect of Context Extraction

- Effectively reduce extracted paths per log entry
- Contextual discrepancies play a crucial role in detecting CPM attacks

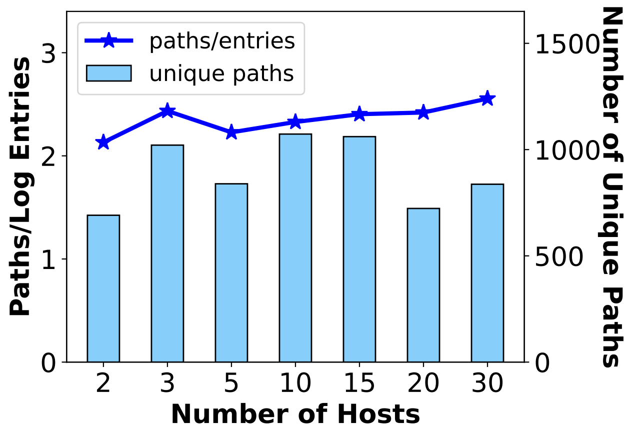

- The number of unique paths stable regardless of network scales and log volumes

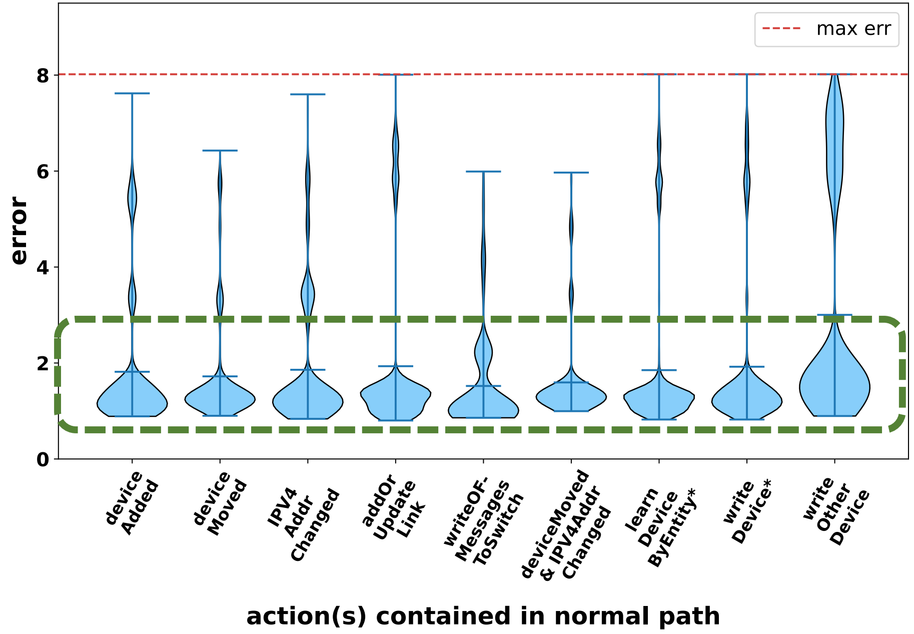

- Isolated actions cannot provide semantic differentiation

Workload & Overhead

- Reduce workloads of manual investigation with acceptable latency and storage overheads
  - Only 6.02% of edges require manual review

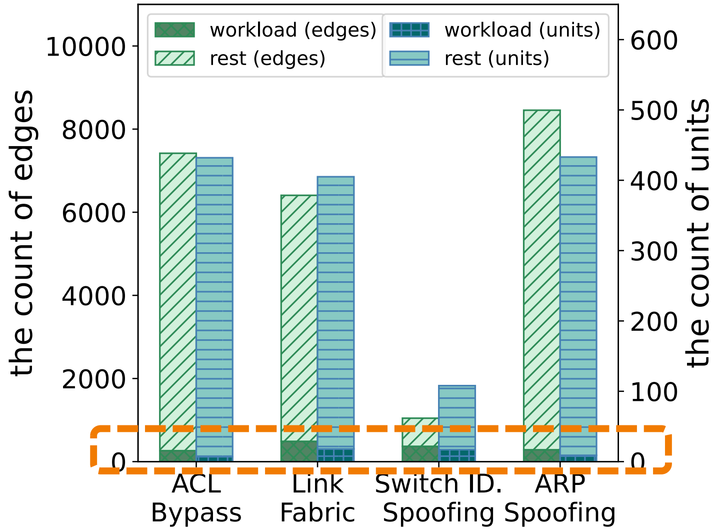

Workload & Overhead

- Reduce workloads of manual investigation with acceptable latency and storage overheads
  - Only 6.02% of edges require manual review
  - RTT extensions average between 1.8% ~ 24% over the uninstrumented controller
  - Audit log data costs 1.3 GB/hr storage overhead

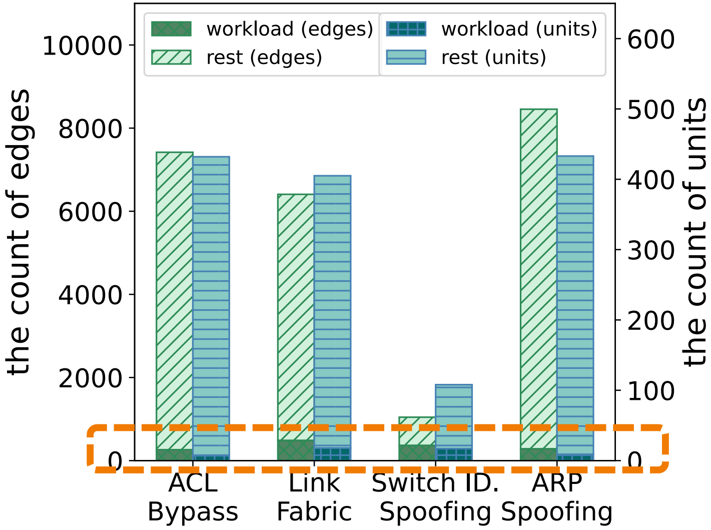

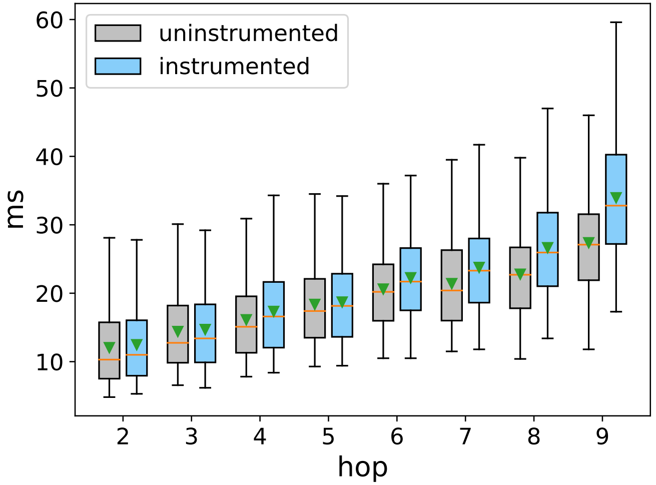

- Each traffic requires the controller to calculate new forwarding rules

Summary

- Our Approach
  - Extracts paths in the provenance graph of SDN controller activities to capture long-term behavior contexts
  - Detects control policy manipulation by identifying deviant contexts based on a prediction model
  - Supports anomaly detection and investigation with minimal reliance on domain-specific knowledge or predefined rules
- Insight
  - Provenance graph contains causal contexts behind the controller's decision-making

ProvGuard: Detecting SDN Control Policy Manipulation via Contextual Semantics of Provenance Graphs

Thank you!

liuziwen@buaa.edu.cn