

Leo: Online ML-based Anomaly Detection at Multi-Terabit Line Rate

Syed Usman Jafri, Sanjay Rao, Vishal Shrivastav and Mohit Tawarmalani

Purdue University

Motivation

Context: ML-based anomaly detection

- Intrusion detection and prevention
- Application and IoT device classification


Why machine learning?

- Captures behavioral and statistical patterns
- Learn complex patterns w/o payload inspection
→ Useful for encrypted traffic


Why in-network anomaly detection?

- Today's networks: Tbps / 100s of Gbps
- Servers can only handle few Gbps
- Programmable switches offer new opportunities
→ Multi-terabit execution of user programs

Traditionally:



Our goal:



Challenges with programmable switches

- Not run-time programmable
 - Program changes → Reboot (downtime)
- No support for mul./div. and floating points
- Limited computation and memory resources

Why decision trees?

- Good match for programmable switches
- No multiplication/division required
- Easily interpretable and high accuracy

Limitations of prior work:

	Runtime Prog.	Not limited by tree dependency	Implementable in ASIC switch	Low ALU usage	Low memory usage
Infocom			✓		✓
pForest	✓			✓	✓
SwitchTree	✓			✓	✓
Illy	✓	✓	✓	✓	✓
Leo	✓	✓	✓	✓	✓

Limitations of prior work

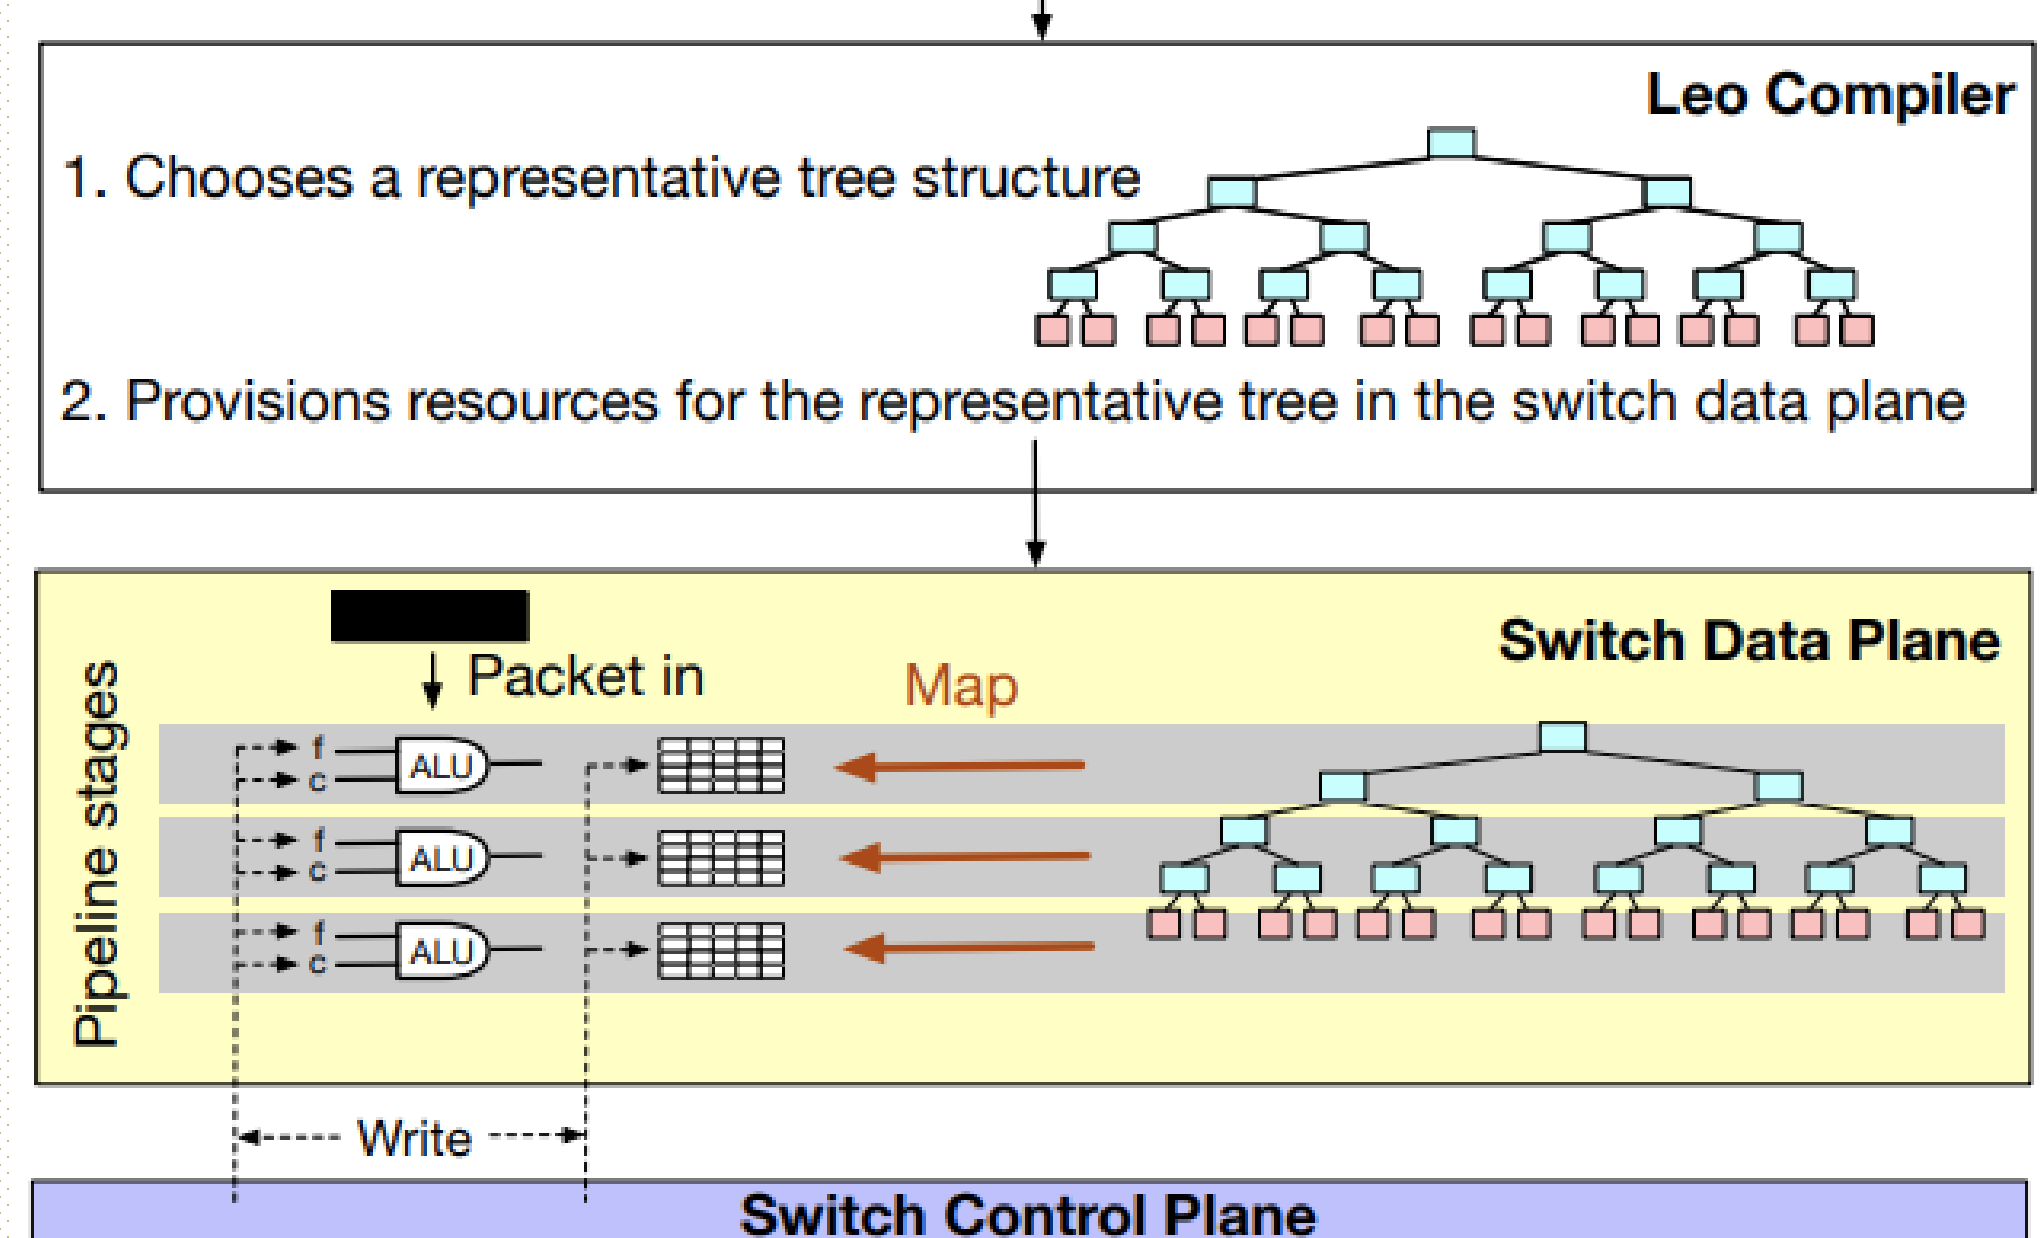
- Not run-time programmable
- Limited tree depth
- Expensive memory requirement

Our solution: Leo

- Support a **class of decision trees**:

- ✓ Run-time programmable
- ✓ Resource efficient
- ✓ Large models supported
- ✓ Multi tera-bit line rate

($D = 4$, $L = 16$, $F = \{f1, f2, f3, f4, f5, f6, f7, f8, f9, f10, f11, f12, f13, f14, 15\}$)



Design

Enabling run-time programmability

- **Multiplexed ALU**: different (feature, constraint) pairs populated to the same ALU

Resource efficient mapping

- **Insight**: a packet accesses one node per tree level
→ Only a single ALU per tree level required

Scaling to larger, deeper trees

- We develop **Sub-tree multiplexing**
- Flattens portions of the tree (sub-trees) to execute in parallel

Evaluation

- We evaluate Leo in a hardware testbed
- Using SRAM, Leo can support classes of trees 2x deeper than the state of the art achieving F1 scores of 0.94 on a IDS dataset

