Speeding Up Network Intrusion Detection

João Romeiras Amado, Salvatore Signorello,
Miguel Pupo Correia, Fernando Ramos

Instituto Superior Técnico, Universidade de Lisboa
Faculdade de Ciências, Universidade de Lisboa

Motivation

- Increasing sophistication of recent attacks
- Need for fast attack detection
- Quality of measurement data

Short-lived network attacks are becoming increasingly common, while existing solutions often take several minutes to perform detection.

R. Miao, R. Potharaju, M. Yu, and N. Jain, "The dark menace: Characterizing network-based attacks in the cloud," in *Proceedings of the 2015 Internet Measurement Conference*, ser. IMC '15, 2015.

M. Moshref, M. Yu, R. Govindan, and A. Vahdat, "Trumpet: Timely and precise triggers in data centers," in *Proceedings of the 2016 ACM SIGCOMM Conference*, ser. SIGCOMM '16, 2016.

Packet sampling's coarse-grained view of the network reduces the effectiveness of intrusion detection.

Sampling introduces a fundamental bias, resulting in degraded performance.

Daniela Brauckhoff, Bernhard Tellenbach, Arno Wagner, Martin May, and Anukool Lakhina. 2006. Impact of packet sampling on anomaly detection metrics. In Proceedings of the 6th ACM SIGCOMM conference on Internet measurement. 159–164.

Anna Sperotto, Gregor Schaffrath, Ramin Sadre, Cristian Morariu, Aiko Pras, and Burkhard Stiller. 2010. An overview of IP flow-based intrusion detection. IEEE communications surveys & tutorials 12, 3 (2010), 343–356.

J. Mai, C.-N. Chuah, A. Sridharan, T. Ye, and H. Zang, "Is sampled data sufficient for anomaly detection?" in Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement, ser. IMC '06, 2006.

Switch-**P**owered **I**ntrusion **D**etection

- Intrusion detection framework powered by **programmable switches**
- **Push-based measurement** approach, reconfigurable at runtime
- **Machine Learning-based** traffic analysis
- Focus on **fast attack detection**

System Design and Architecture

Switch-Powered Intrusion Detection

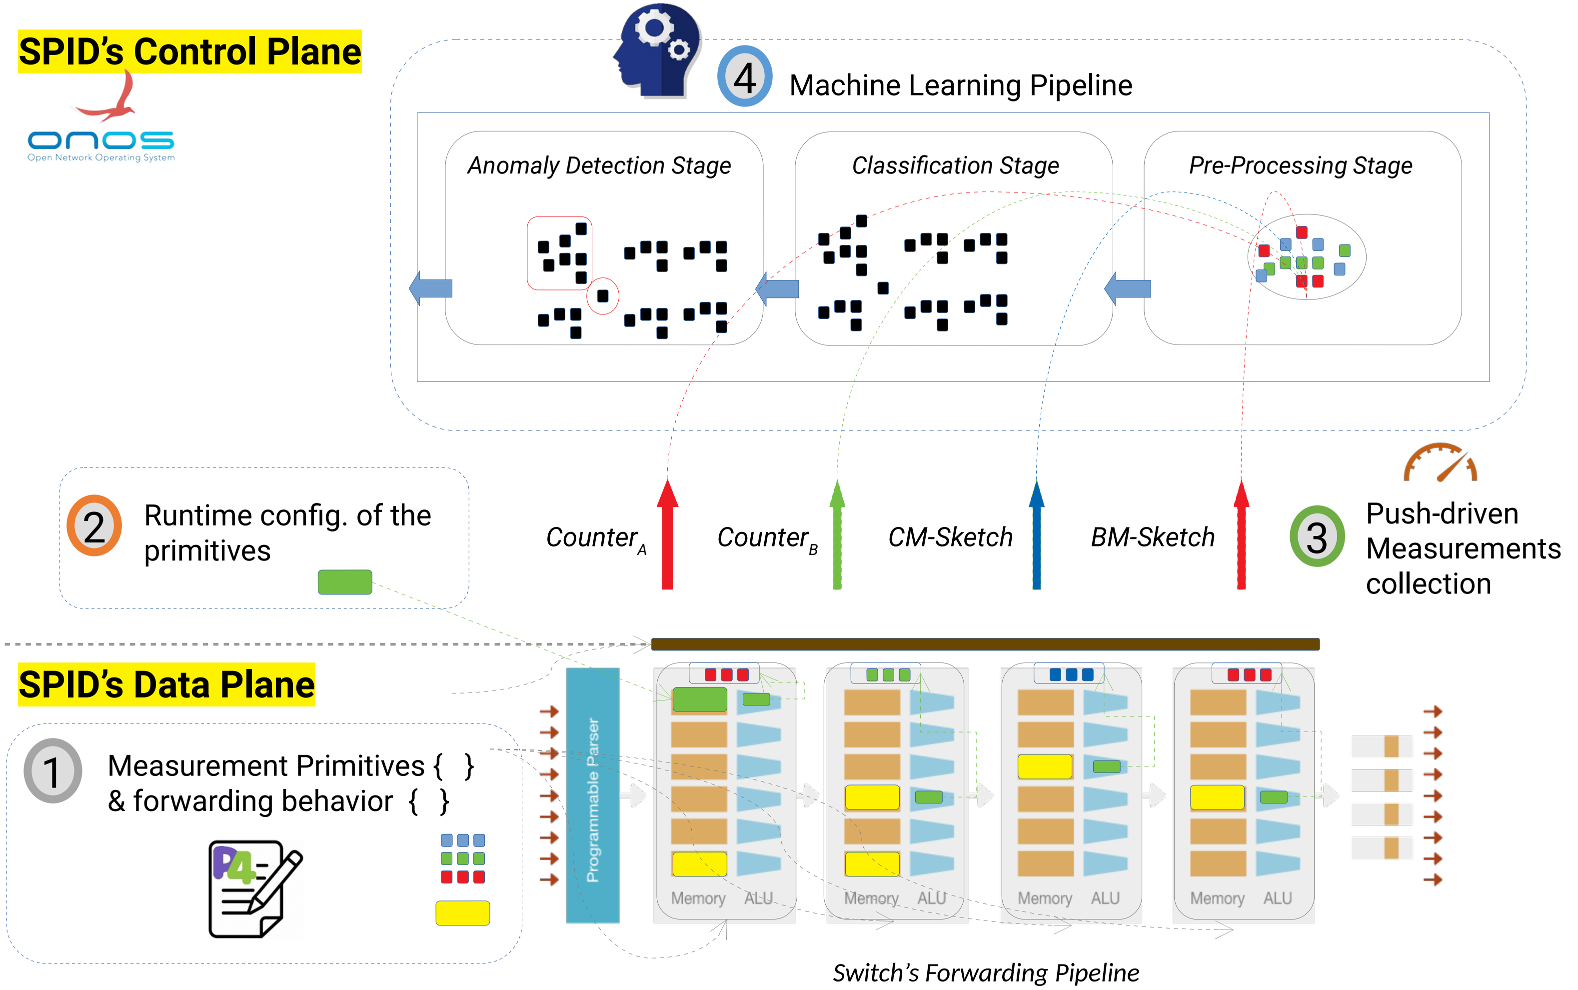

Machine learning-based traffic analysis.

Push-based switch-driven statistics collection.

Rich set of **packet summaries** stored in the switches, reconfigurable at runtime.

Measurement Primitives

Flow Statistics	**Sketching Algorithms**
Number of packets/bytes	Count-min
Source/Destination IP	Bitmap
IP Protocol	AMS
Source/Destination Ports	K-ary
TCP Flags	MV-Sketch
ICMP Type/Code	HyperLogLog
(...)	(...)

The operator is able to reconfigure the active counters on all switches at runtime.

Each switch can be optimized for different monitoring purposes.

Measurement Primitives and Runtime Config.

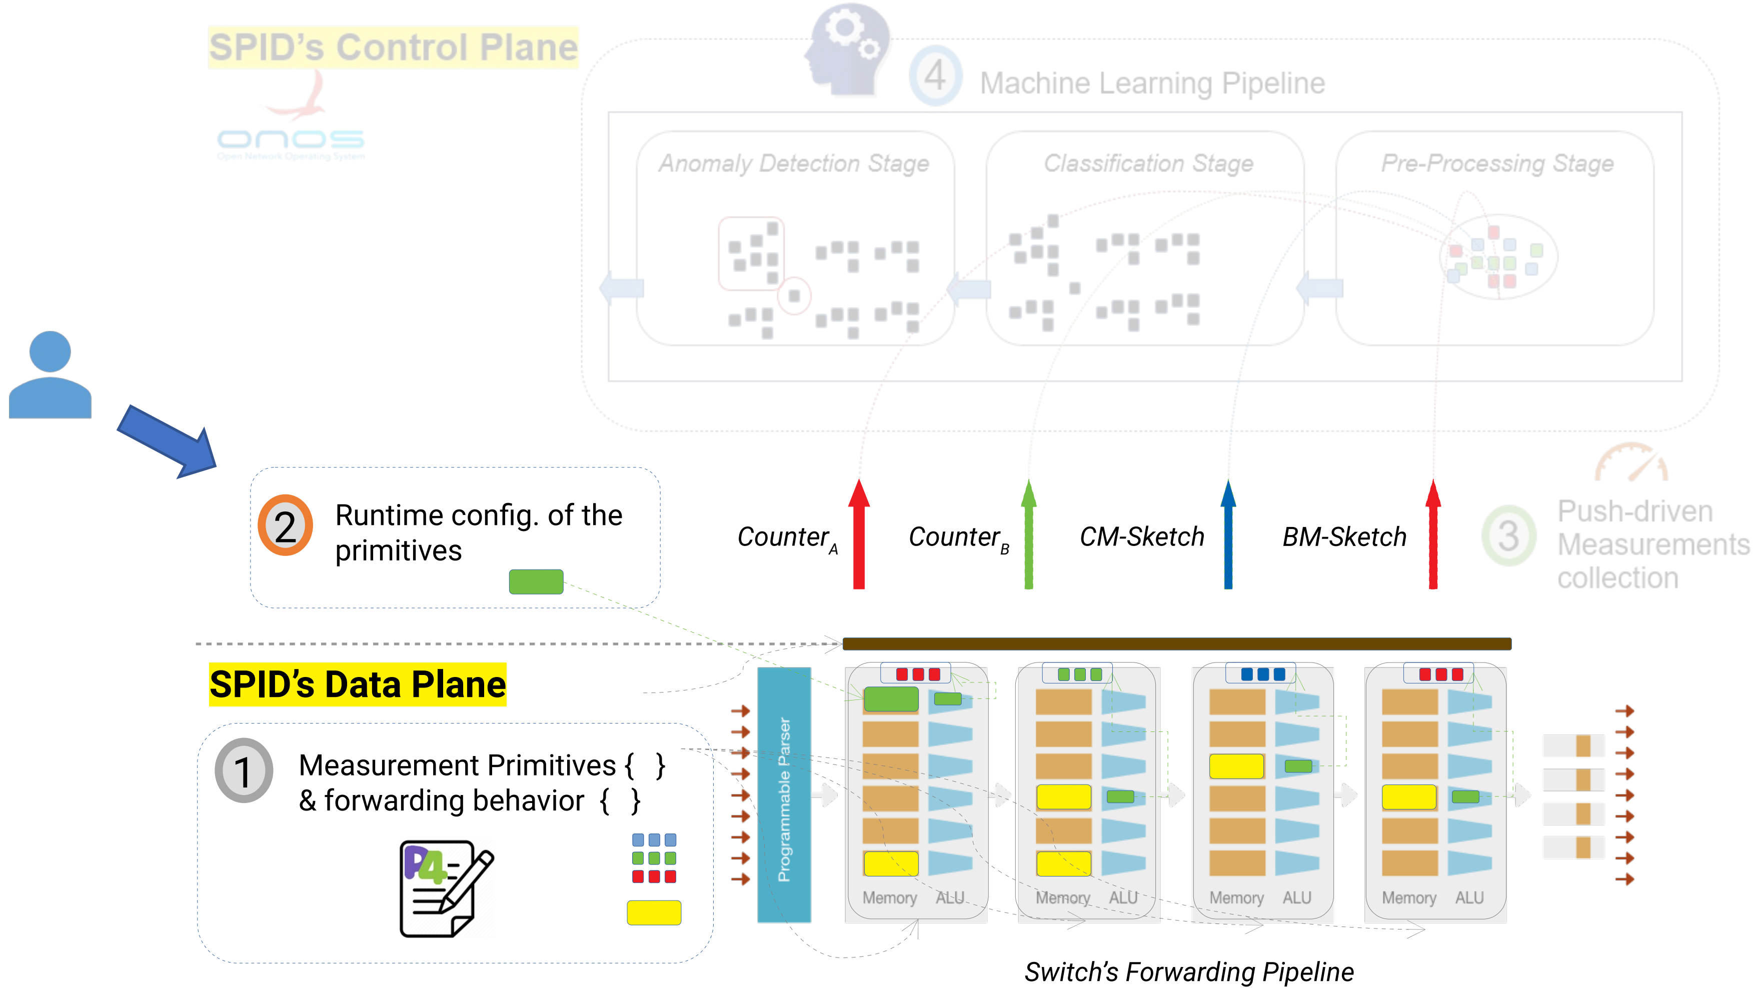

Each switch's available memory is **dynamically allocated** between all active counters.

The operator is able to **reset all counters** during runtime.

Push-driven Measurement Collection

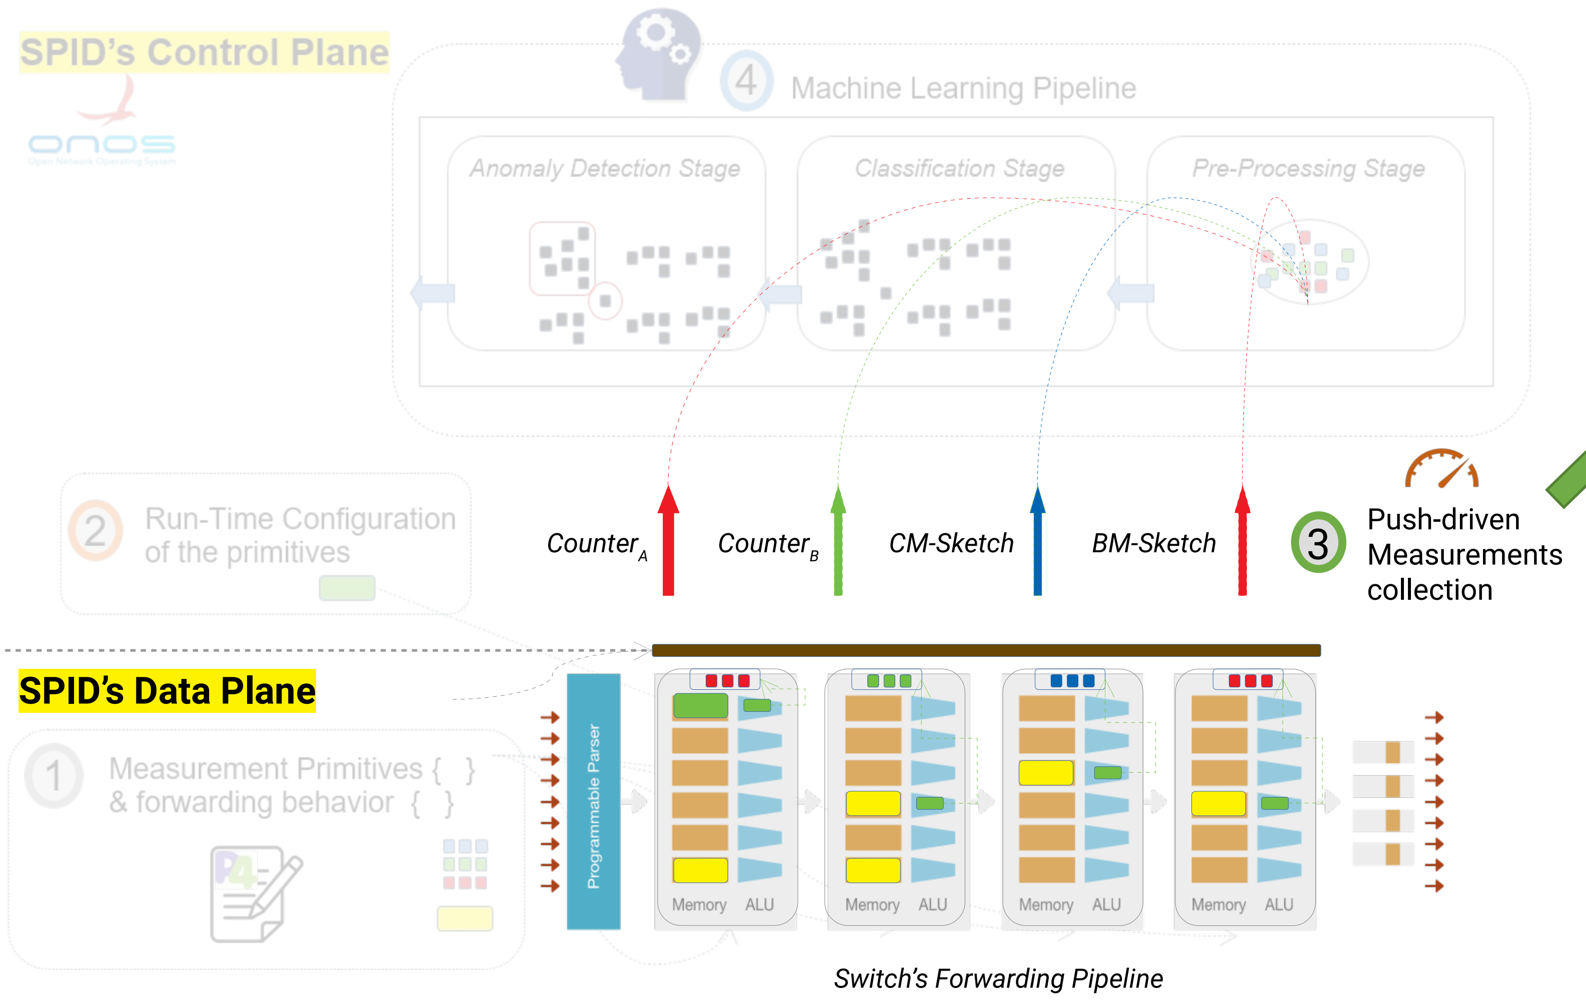

Traffic change detection sketches will serve as triggers for the push-based collection.

Faster alerts

Switch Proactiveness

Relieves the control plane of the burden of performing polling actions.

Machine Learning Pipeline

Goal: Perform flow aggregation according to their characteristics, aiming to detect potential anomalies in the form of outliers.

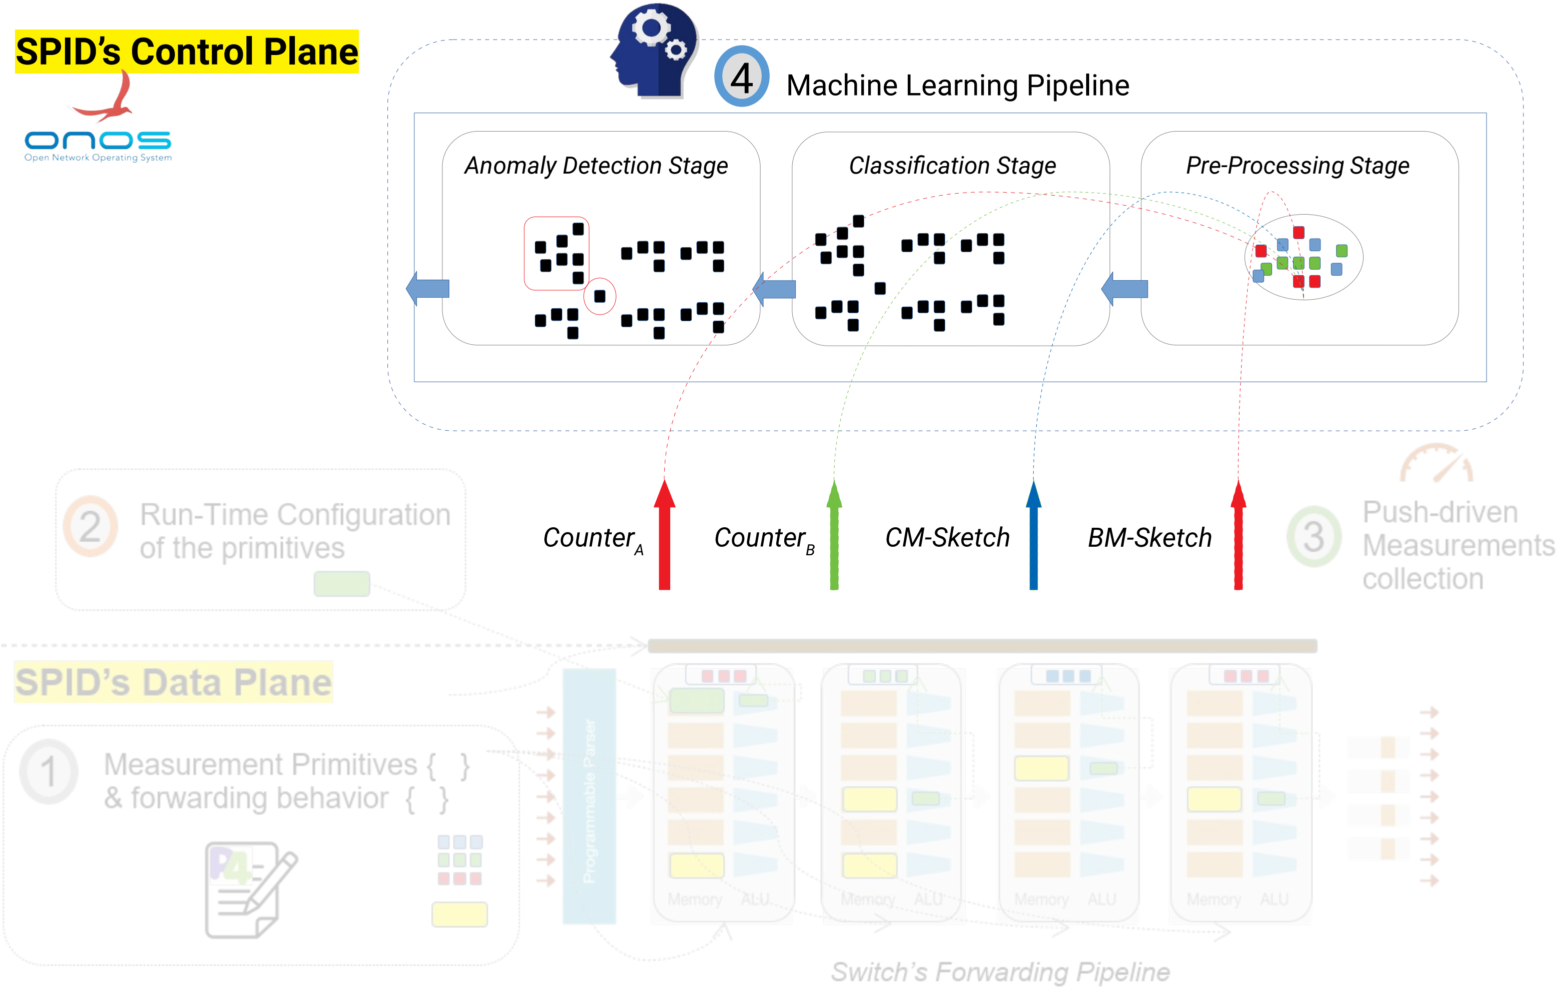

The pipeline is immediately executed when a trigger event is received from the data plane.

SPID's collection of multiple measurement primitives is essential to increase the number and variety of network features available as input to the detection system.

Preliminary Evaluation

Preliminary Evaluation

- Detection of unknown attacks
- Stream-based over sample-based
- Detection time

The evaluation was performed with **real traffic datasets** containing multiple labeled attack instances.

While SPID observes all packets, we also tested a sample-based approach that performed a sample of **1/500 packets**.

Evaluation: Detection of Unknown Attacks

Attack Type	Solution	TP	FP	Precision	Recall
TCP SYN Flood	SPID	40.0%	66.0%	37.7%	99.0%
	CM Sketch	30.0%	69.7%	30.1%	98.6%
	Sampling	0.0%	100%	0.0%	0.0%
Ping-of-Death	SPID	93.3%	44.8%	67.5%	99.9%
	CM Sketch	30.0%	94.2%	24.2%	98.2%
	Sampling	46.7%	68.4%	40.6%	94.4%

Across tested attacks, SPID always has a higher precision percentage than the other baseline NIDS.

A combination of **multiple measurement primitives** is much better than any single metric.

Very preliminary results with a basic ML approach!

Evaluation: Detection Time

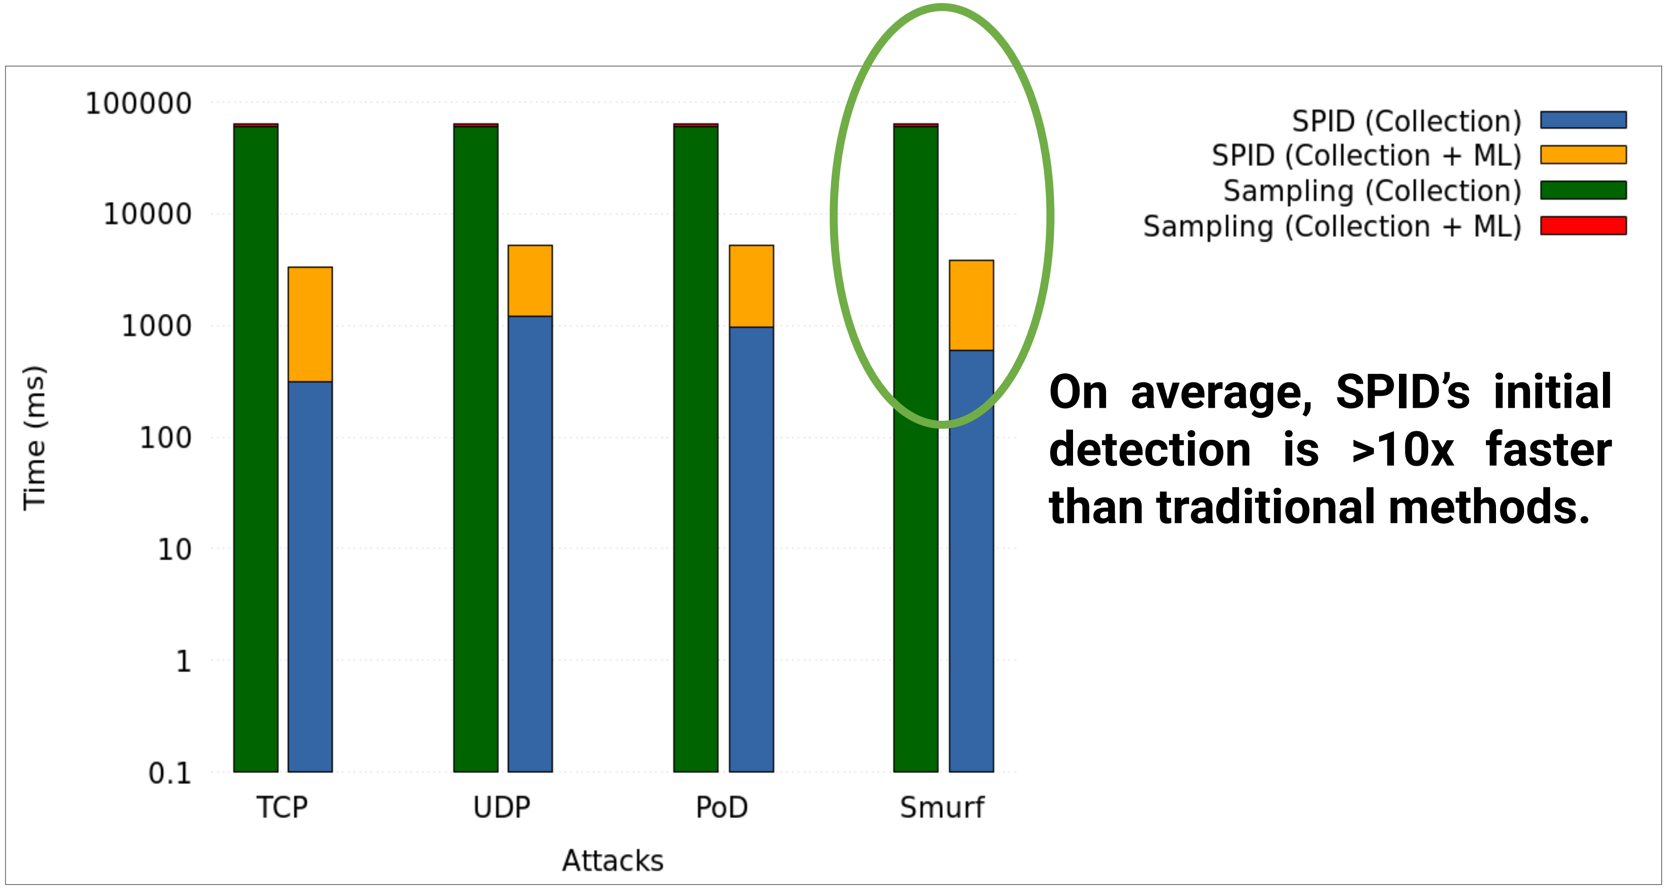

On average, SPID's initial detection is >10x faster than traditional methods.

Sampling: The detection time of a sampling-based approach is inherently constrained by the sampling frequency.

SPID: A push-driven approach detects anomalous patterns as soon as they emerge in the data plane.

Current Status

Our preliminary experiments offer confidence on the potential of programmable switches in improving network-based IDSs, namely given:

a) The ability to collect and reconfigure during runtime a diversity of different measurements at the switch-level, including sketching algorithms, points towards an improvement in detection precision

b) Potential of a **push-driven approach** to speed up intrusion detection

c) Use of anomaly detection techniques to filter alerts from the data plane, allowing the operator to focus only on the more relevant traffic statistics

Future Work

- Design and implementation of additional (and refinement of existing) **measurement primitives** in P4, along with **lightweight traffic change detection algorithms** to enable better data plane triggers

- Deployment and testing of SPID on **P4-programmable hardware**

- Explore modern **anomaly detection approaches** to improve the precision of SPID to the level required by intrusion detection environments