VOZES DIVERSAS
DIFERENTES SABERES

Salão UFRGS 2018 — Conhecimento Formação Inovação

SALÃO DE INICIAÇÃO CIENTÍFICA XXX SIC

15 A 19 OUTUBRO CAMPUS DO VALE

UFRGS PROPESQ

Meta-Data Management on Programmable Data Planes

Lucas Castanheira and Alberto Schaeffer-Filho

Institute of Informatics, Federal University of Rio Grande do Sul, Porto Alegre, Brazil

Introduction

- Looking inside a network in order to detect suspicious behavior from an end user is a difficult task (the amount of data is overwhelming).
- Monitoring what happens in Software-Defined Networking (SDN) is usually the job of the controller and involves a trade-off between high fidelity (gathering all data we can to better understand things) and low impact (making our data collection cause as little stress to the network as possible).
- With programmable data planes we are able to employ all the distributed hardware of data plane routers in monitoring the network to look for suspicious behaviors.
- However, since the data plane operates on a tight schedule, we cannot run most detection solutions (techniques which follow up monitoring, such as machine learning) on the routers directly.
- We created a system that allows us to perform monitoring on the data plane and efficiently shift the data to the control plane for extensive processing.
- Our system, is composed of two subsystems: Monitoring and Gathering.

Gathering

- Distributively consolidates all the information about the suspect flow that is inside routers and stores it in the controller.
- Separates the network into clusters.
- Tries to minimize the impact on the controller of requesting data to multiple routers.
- Controller injects a Crawler Packet (CP) that makes routers talk to each other and pass the request along, only sending the resulting data to the controller.

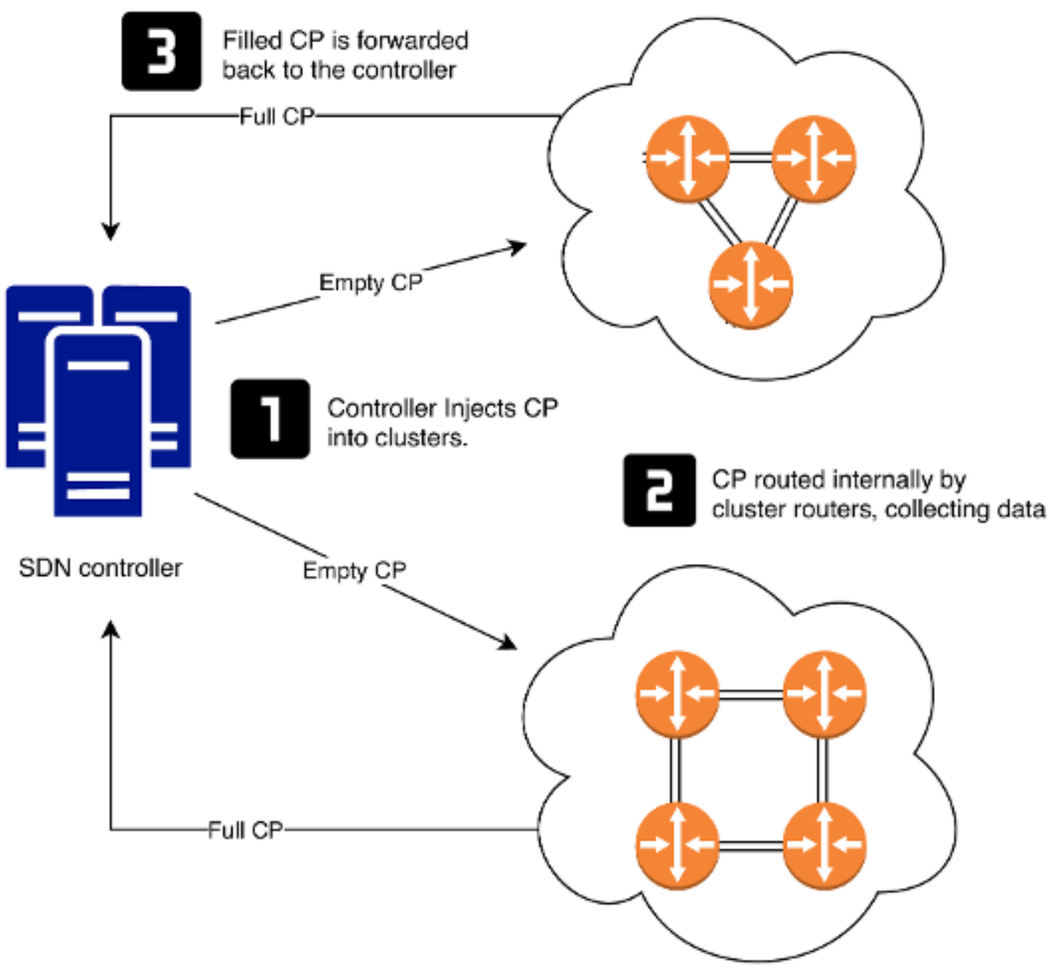

Monitoring

- Looks at every flow and tracks the amount of traffic it generates (lightweight monitoring).
- Whenever a flow starts creating larger amounts of traffic, the system triggers monitoring on a more comprehensive scale for that flow, locally storing all the timestamps and relative sizes for each upcoming packet pertaining to the flow.
- If the flow does not cease after this analysis, the monitoring phase warns the controller about the flow.
- The controller then triggers the gathering system.

Heavyweight Monitoring System

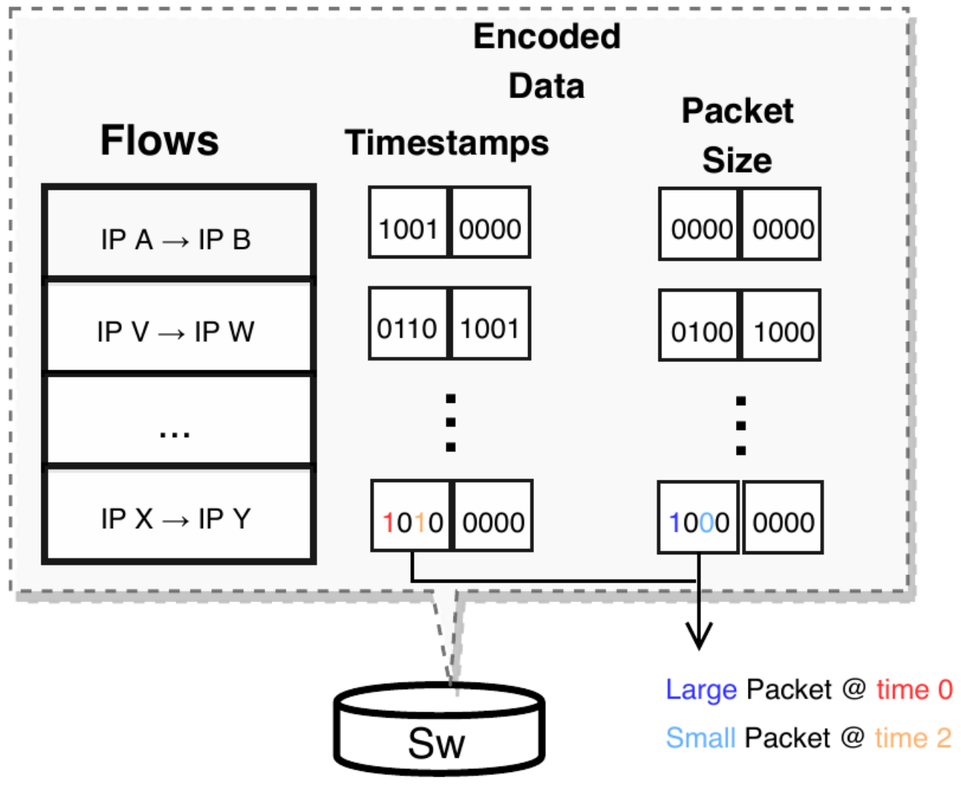

Results

- Preliminary results indicate that we were able to capture all of the desired timestamps and relative packet sizes correctly on the heavy monitoring phase.
- The constant Lightweight monitoring creates a small rise in overall packet latency.

Conclusions and Future Work

- Shifting a part of the monitoring responsibilities to the data plane can have powerful advantages over traditional methods (e.g snapshotting) in both efficiency and fidelity of collected data.
- Our system provides us with meta-data about flows that is naturally fit for building a data-set for Machine Learning detection techniques.
- We aim to continue our work by developing a detection system running on the controller, based on data collected by our data plane system.