CREASE – Building a Network Debugger for FABRIC

ALEXANDER WOLOSEWICZ (IIT), VINOD YEGNESWARAN (SRI), ASHISH GEHANI (SRI), NIK SULTANA (IIT)

TECHNICAL ASSISTANCE FROM KOMAL THAREJA AND MERT CEVIK

Motivation

Custom fork of BMv2 to provide checksums of its state

Provenance recorded in SPADE, annotated flows crossing the switch

For tracking changes in switch state over time

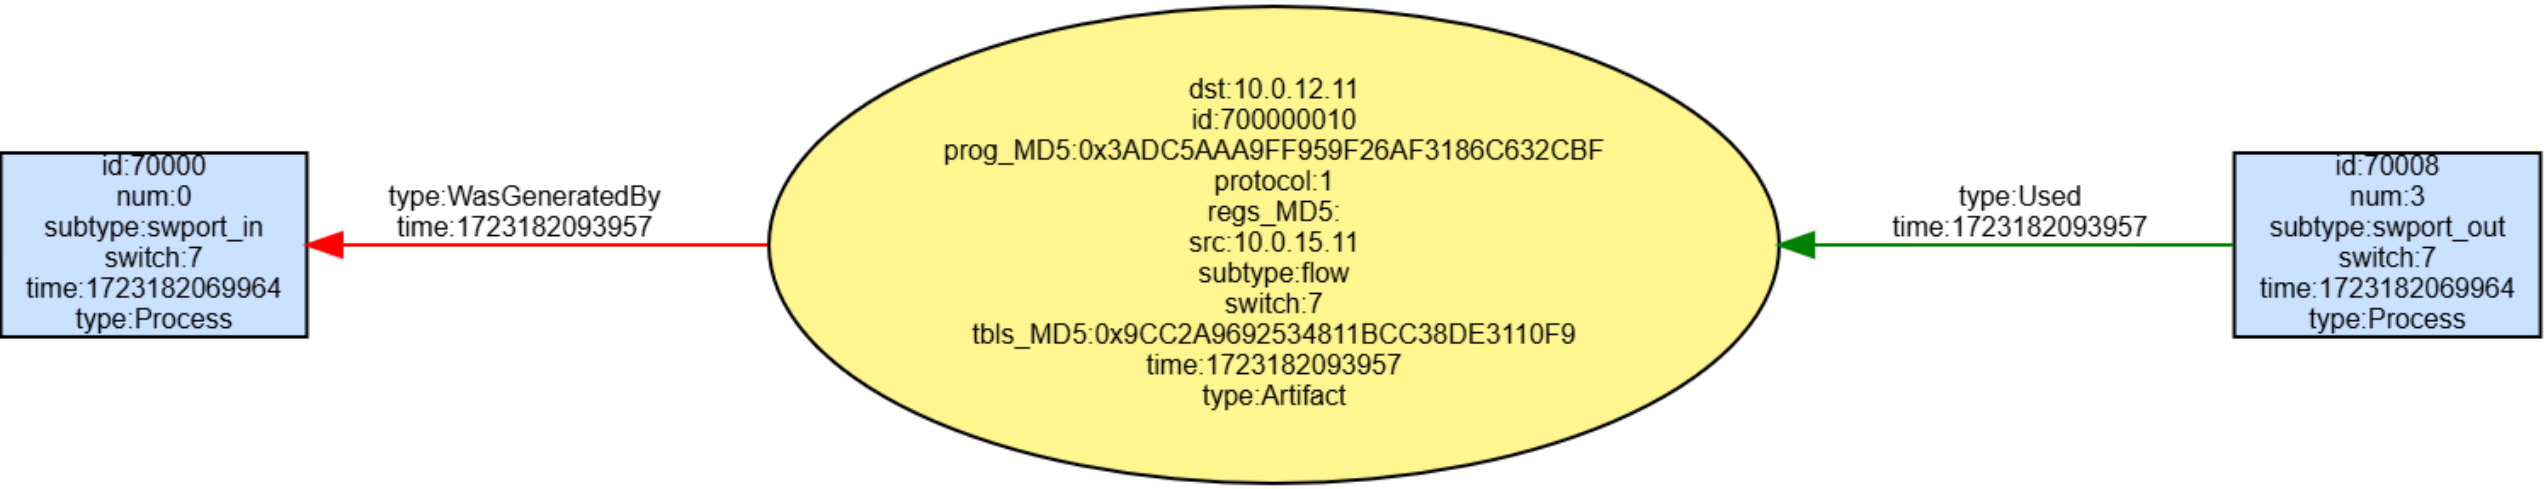

Motivation

Very useful for debugging network issues such as misconfigurations

Query-able, visually makes sense, one place instead of many

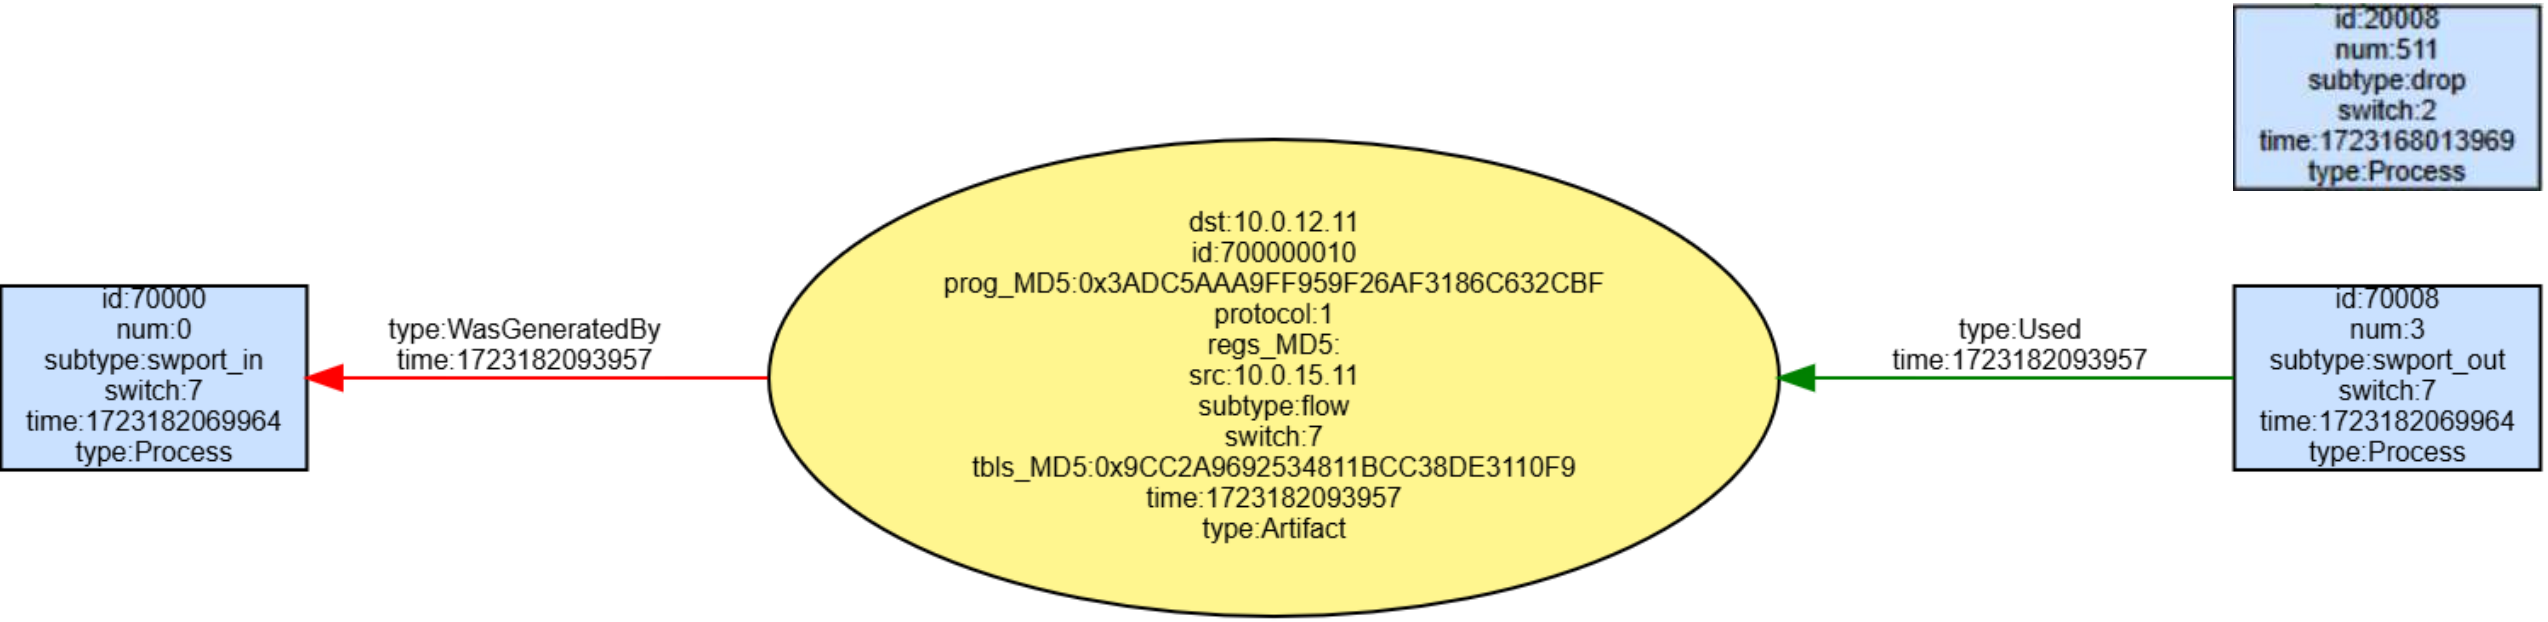

Designing the Debugger

Generalize the BMv2 fork to function independently

FABRIC is a testbed – freely insert VMs

Use these VMs to monitor and manipulate traffic, but abstract them from the experimental topology

- To the user, topology appears to function largely as if the debugger isn't there

These monitors insert a trailer that identifies the packet, enabling tracing the history of a packet across the network

Designing the Debugger

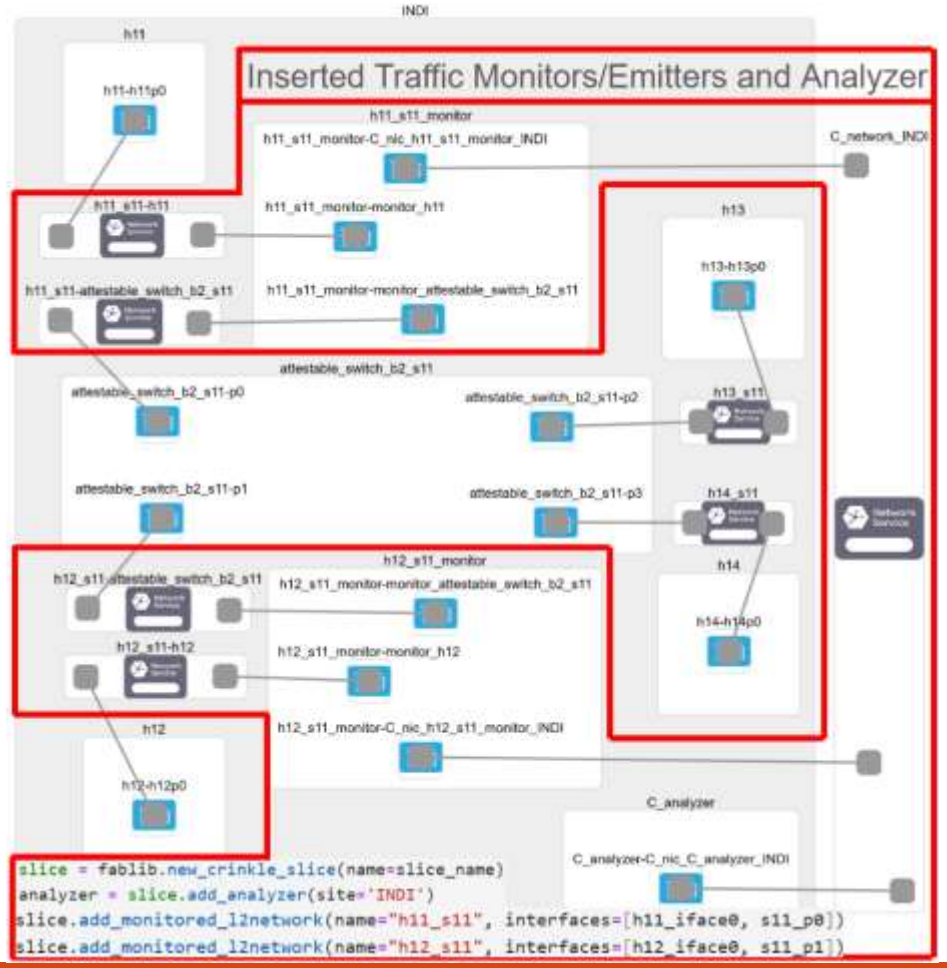

Fork Fablib to keep a familiar and easy-to-use design

Abstract adding and interacting with the monitors – user just interacts with their normal slice, and a simple front-end API to the debugger

Designing the Debugger

Collected packet histories stored in SPADE

Query-able, and produces graph visualizations

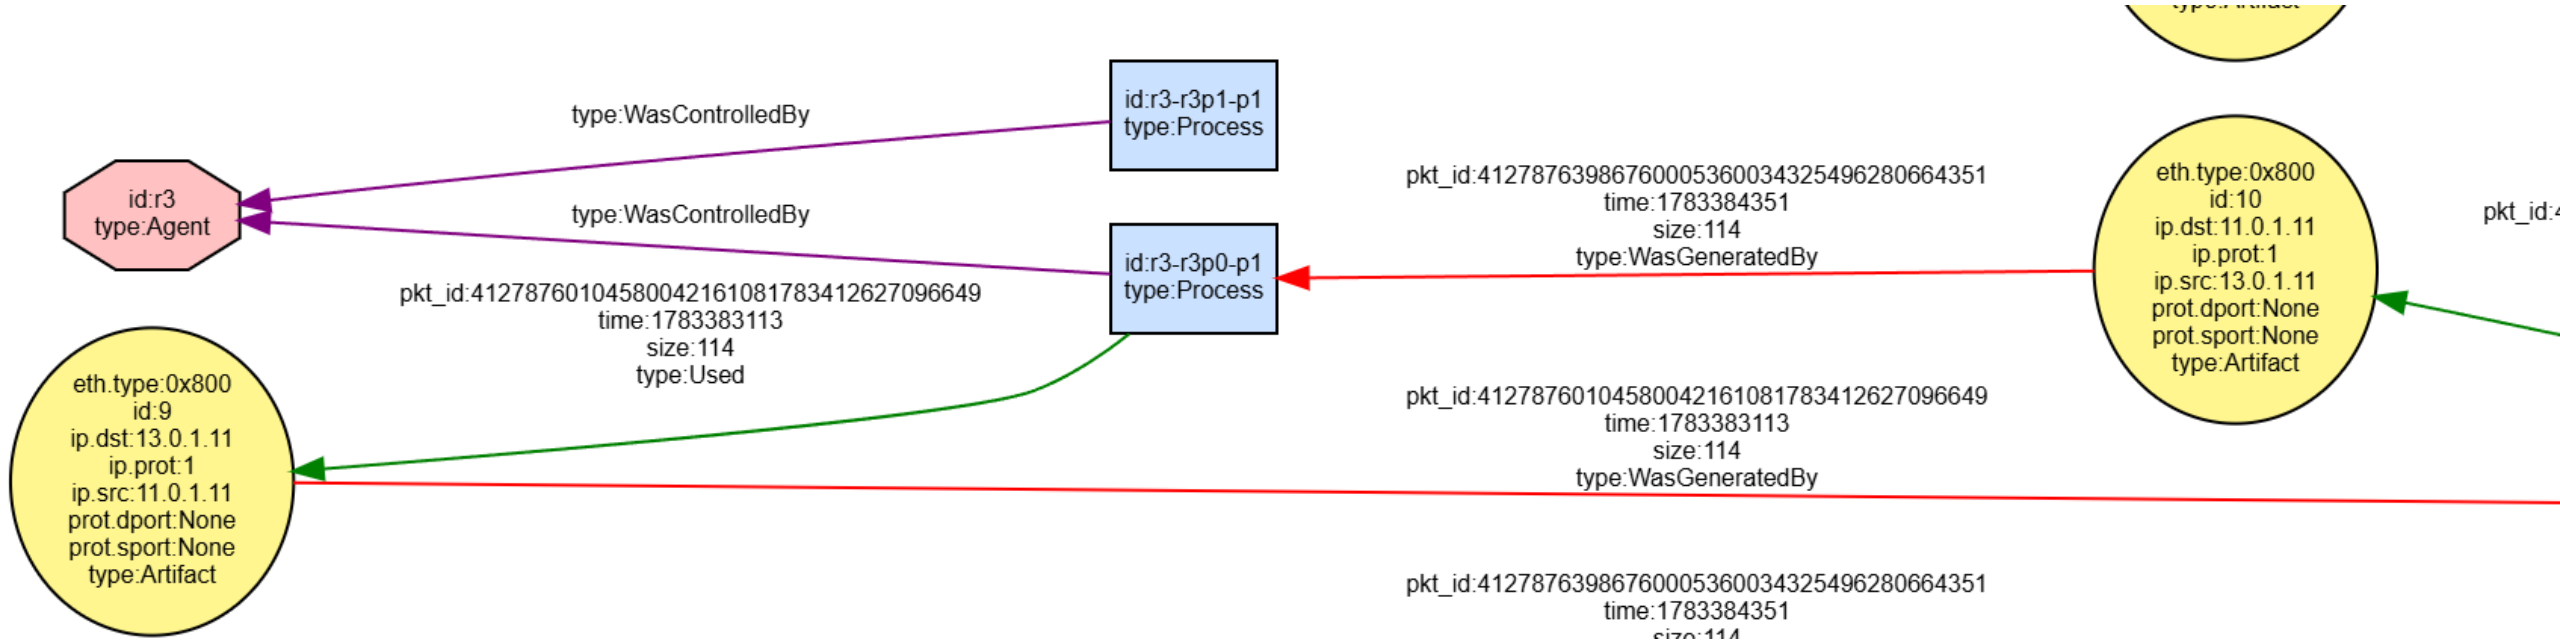

Designing the Debugger

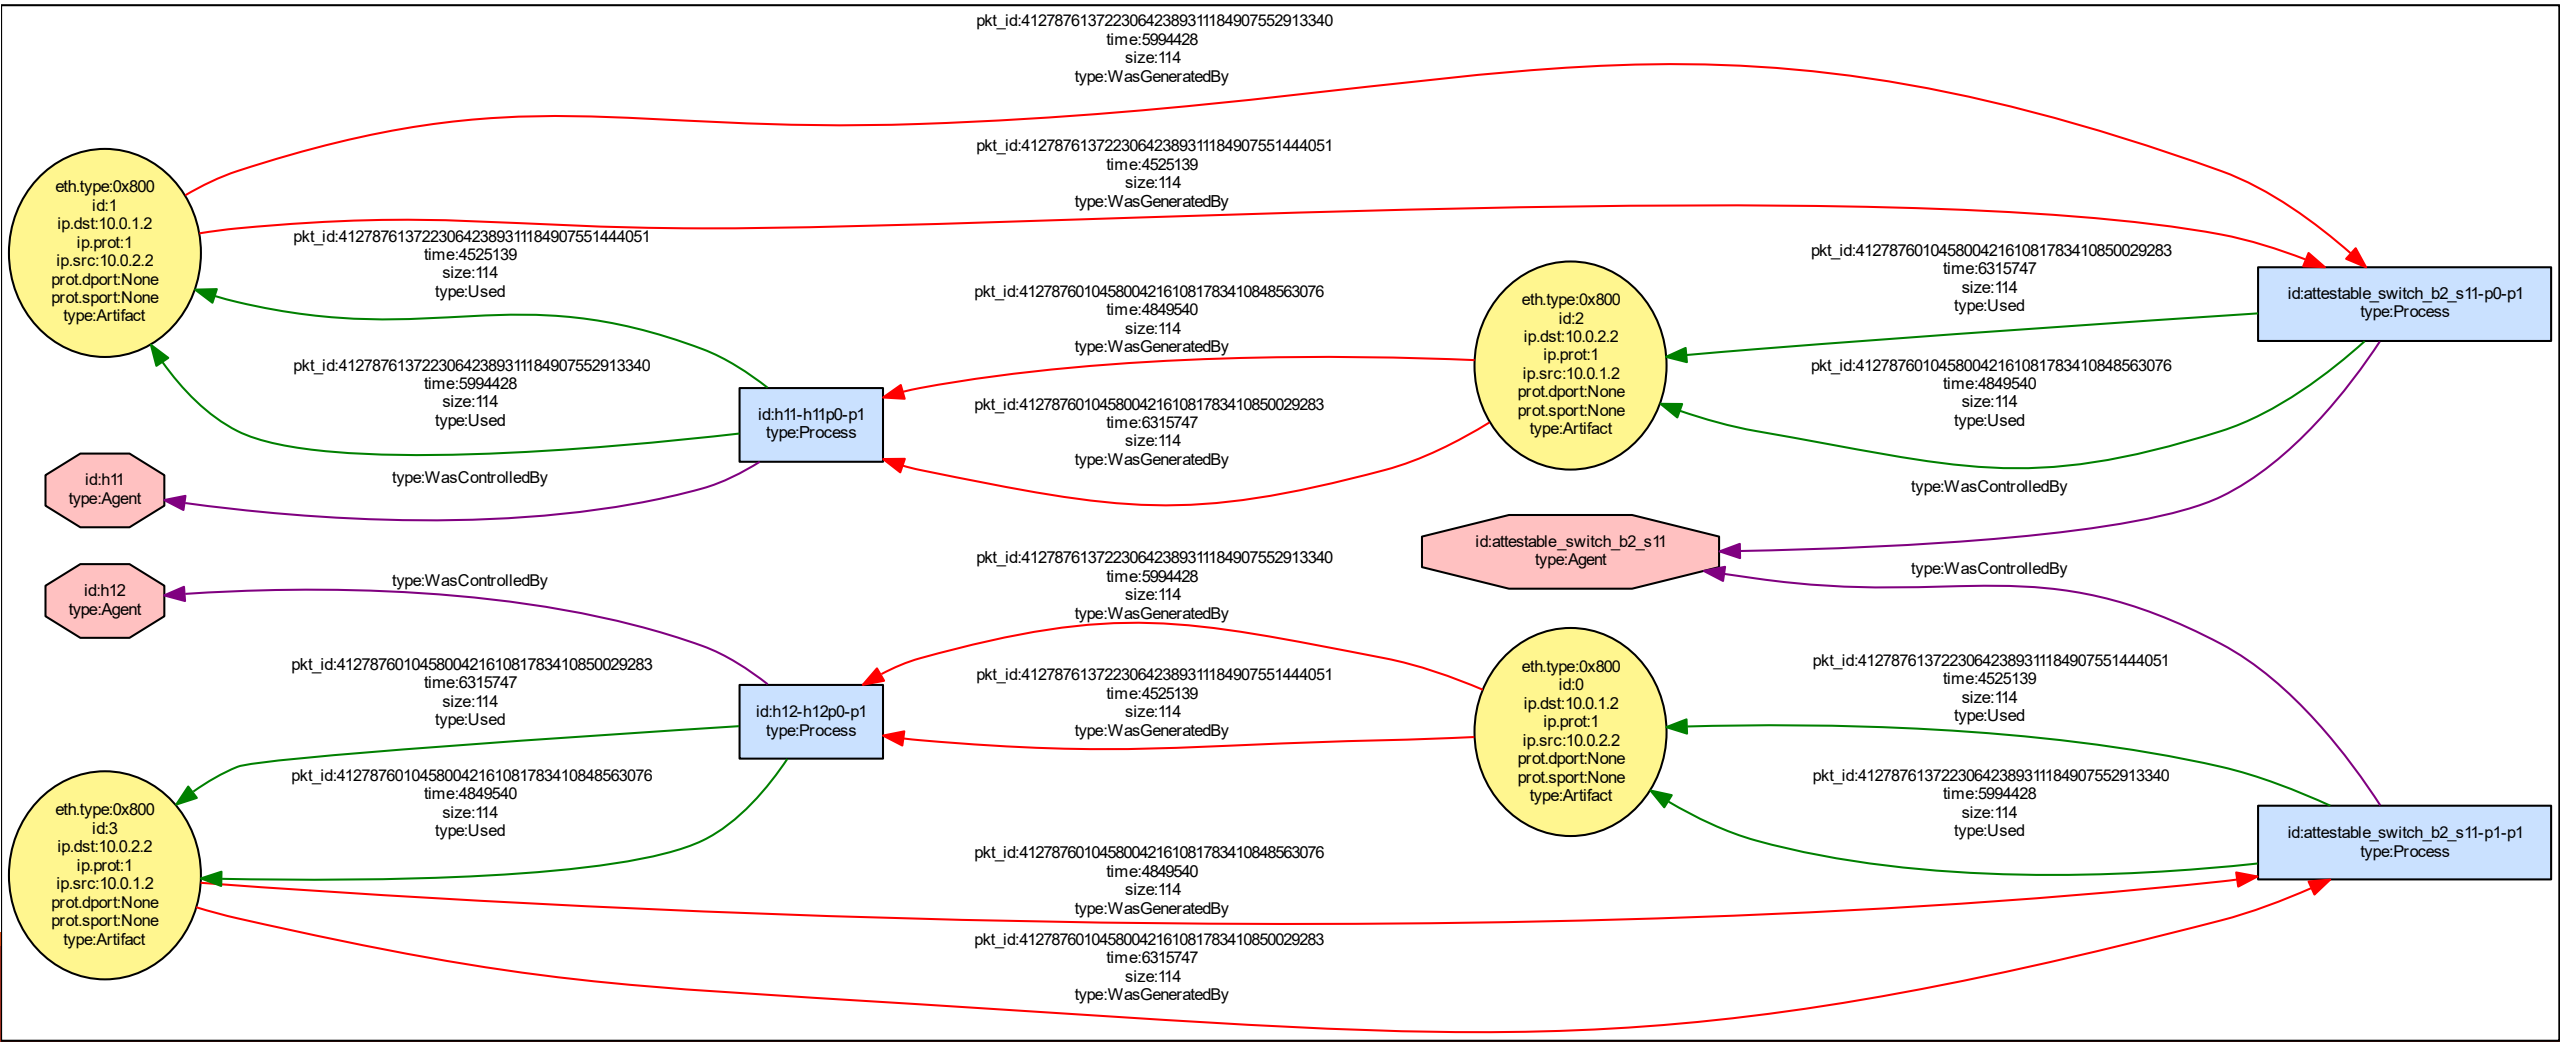

Using the Debugger

Previous images – get_graph()

- Automatically builds and downloads the graph
- Can be filtered by time, tcpdump filters, pkt_id

Using the Debugger

Previous images – get_graph()

- Automatically builds and downloads the graph
- Can be filtered by time, tcpdump filters, pkt_id

Probe()

- Build a custom packet, send it at a target interface
- Gives the graph for that one packet traversing the network

Using the Debugger

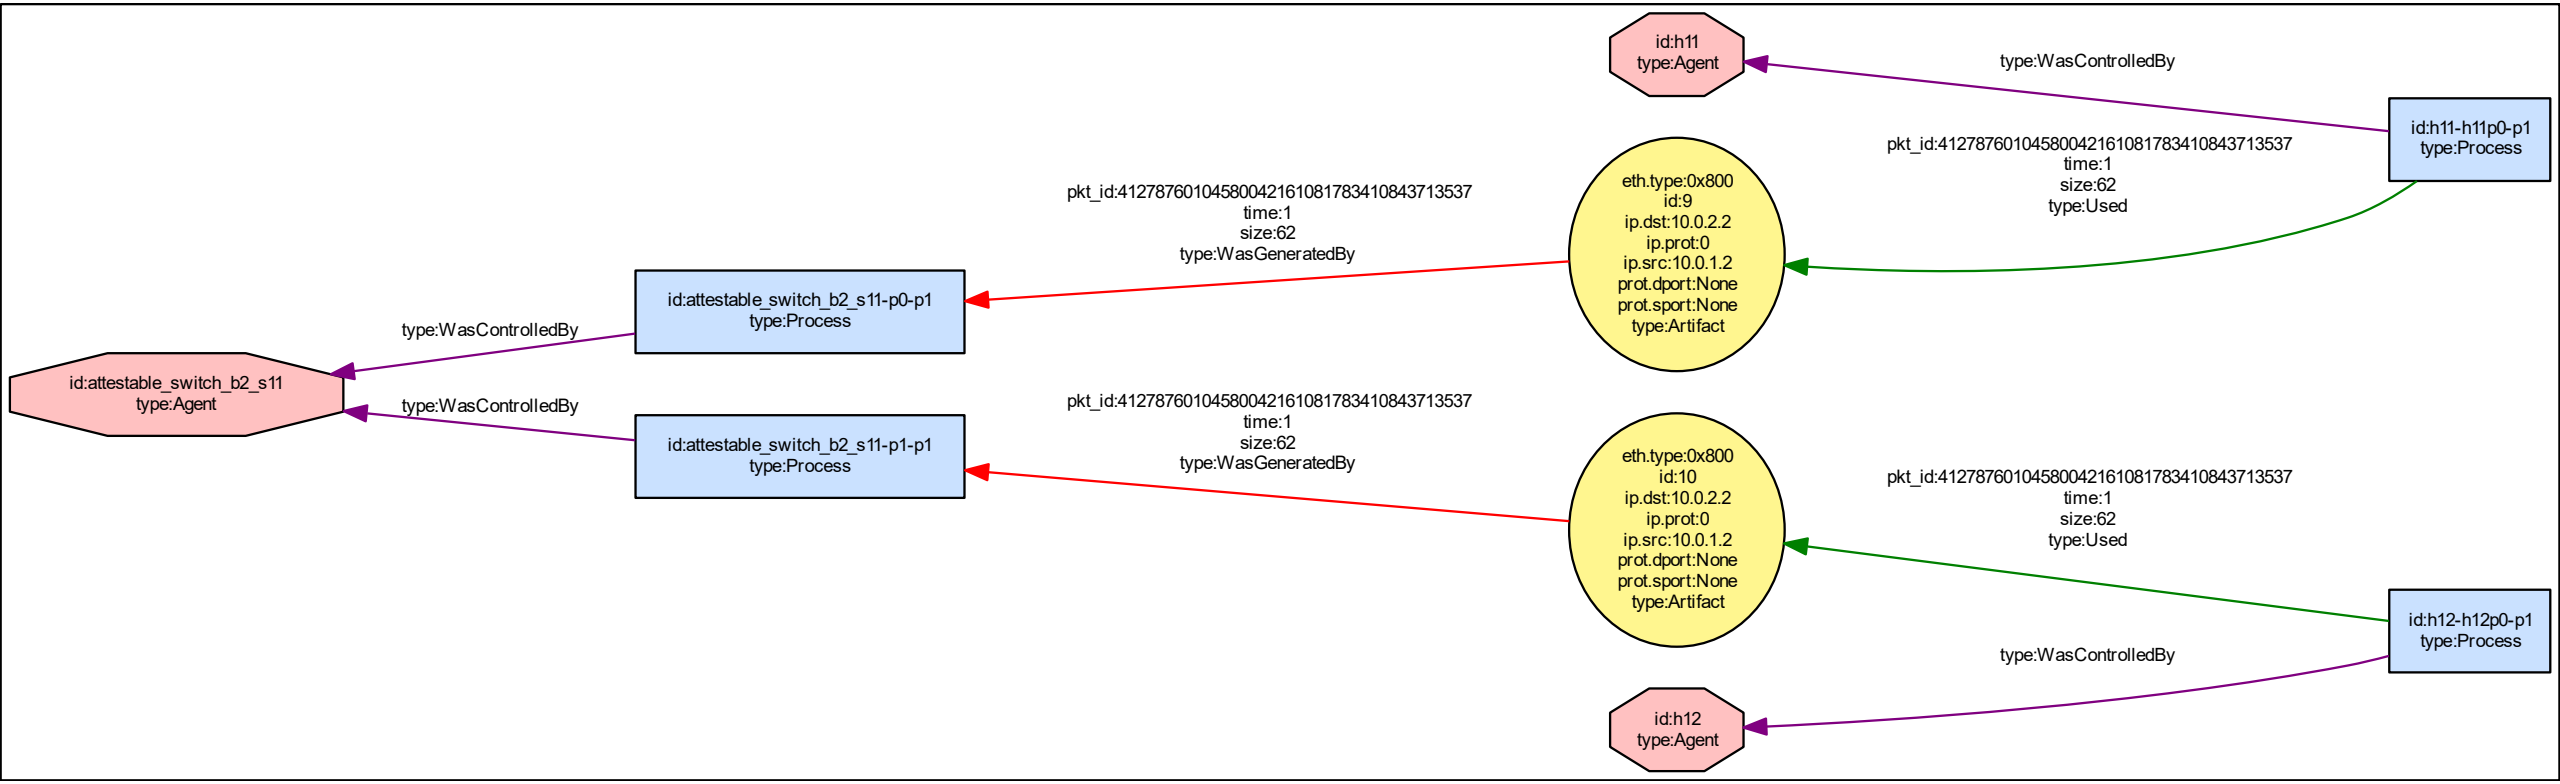

Future Work

DPDK instead of P4-BMv2 driven

More general (user-specified) field support instead of IP

Long-term: automated fault localization

Want to help?

Looking for specific example problems to use for designing and evaluating the debugger

Feedback on the tool itself

Interest in using early versions of the tool

Connect with Nik afterwards

Email me: awolosewicz@hawk.iit.edu