EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

FACULTY OF
SCIENCE
Communication Networks

P4TG: 1 Tb/s Traffic Generation for Ethernet/IP Networks

Steffen Lindner, Marco Häberle, Michael Menth

bw NET

http://kn.inf.uni-tuebingen.de

Motivation

- New protocols & network equipment needs to be tested with realistic traffic rates
- Traffic generators (TGs) used for this purpose

- The top 10 used TGs in the literature are all software-based!
  - iperf2
  - Netperf
  - Moongen
  - …

- 100+ Gbit/s difficult to generate with software
  - Need hardware acceleration
  - Hardware based TGs very expensive ($$$$$)

- Multi-Port (several 100 Gbit/s) testing for business-grade switches/routers not feasible with software TGs

Idea

- Traffic generation with P4 and Intel Tofino™ ASIC (< 8.000€)
- Intel Tofino™ offers built-in capabilities for traffic generation

- We implement measuring functions and configuration in P4 + GUI
- Constant bit-rate & poisson traffic

+

=

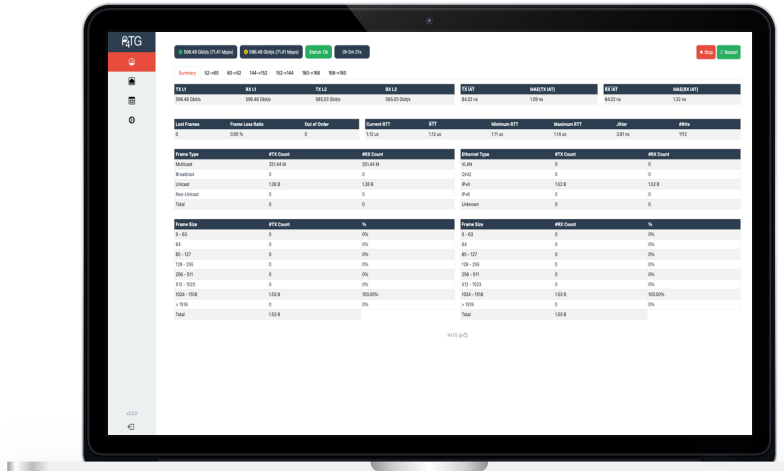

Background (I)

- Intel Tofino™ ASIC
  - 3.2 Tbit/s or 6.2 Tbit/s P4 programmable switching ASIC (Gen. 1)
    - Our Edgecore Wedge supports 32x 100 Gbit/s ports
  - 12.8 Tbit/s P4 programmable switching ASIC with 32x 400 Gbit/s ports (Gen. 2)
  - ~~25.6 Tbit/s P4 programmable switching ASIC with 64x 400 Gbit/s ports (Gen. 3)~~

- Intel Tofino™ ASIC allows for internal traffic generation
  - Up to 8 different packet (byte) descriptions with periodic timer for packet generation

https://www.edge-core.com/productsInfo.php?id=335

Background (II)

- **P4: P**rogramming **p**rotocol-independent **p**acket **p**rocessors
  - High-level programming language to describe data planes
  - Target-specific compiler maps P4 program to hardware

```
control MyPipeline(inout headers hdr, inout metadata meta, inout
standard_metadata_t std_meta) {

  /* Declarations region */
  table ipv4_lpm { … }
  action ipv4_forward(…) { … }
  …

  apply {
    /* Control Flow */
    if(hdr.ipv4.isValid()){
       ipv4_lpm.apply();
    }
  }
}
```

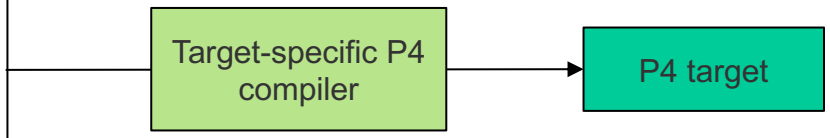

  - P4 defines low level (packet processing) operations

⇒ Fully programmable data plane

  - Limited only by expressiveness and features of P4 (and not by vendor)

Background (III)

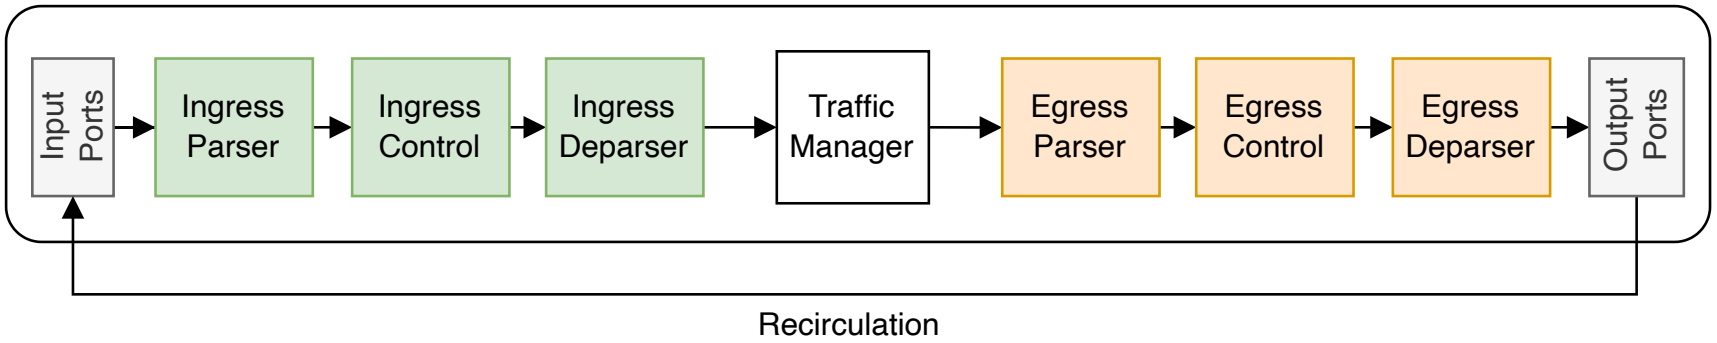

- P4-programmable
  - Ingress/Egress Parser
  - Ingress/Egress Control
  - Ingress/Egress Deparser

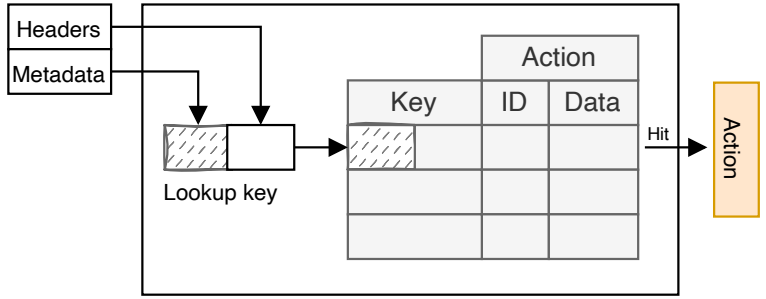

Match+action table used in ingress/egress control

P4TG: Overview

- Leverage internal traffic generator for packet generation
- Packet header rewrite for traffic randomization
- VLAN & Q-in-Q encapsulation support
- Up to 10x 100 Gbit/s traffic generation

- Measure several metrics directly in the data plane (P4) for highest precision
  - L1/L2 TX & RX rates
  - Per stream TX & RX rates
  - TX & RX frame sizes and types (unicast, multicast, broadcast, IPv4, IPv6, VLAN, Q-in-Q)
  - Packet loss, out of order
  - TX & RX inter-arrival times (mean and mean-absolute-error)
  - Round-trip-time (RTT; sampled)

Concept (I)

- Each of the 10x P4TG ports is associated with two recirculation ports
  - P_{out}^{R} and P_{in}^{R}

- Packets received on port P_{in} are recirculated to port P_{in}^{R}
- Packets sent on port P_{out} are initially sent to port P_{out}^{R}

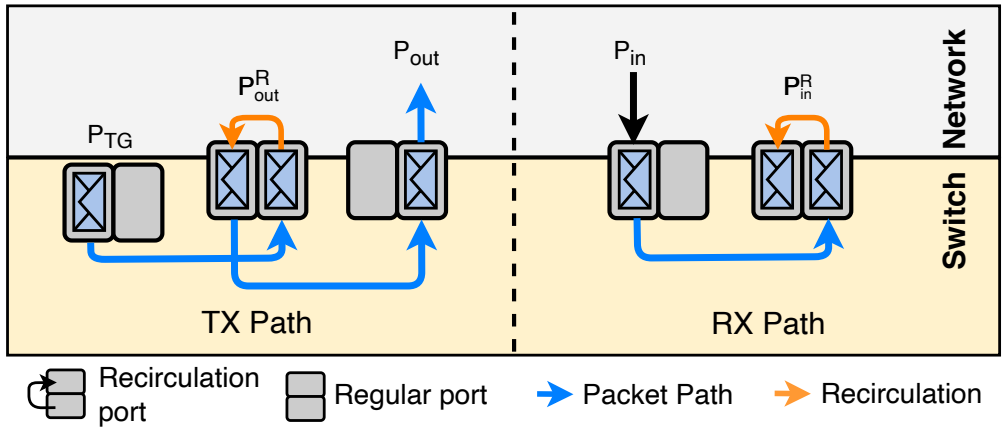

- Statistics are collected during recirculation

Concept (II)

- Measurements
  - 64-bit registers to store total TX & RX bytes per port
    - Hardware timestamps with nanosecond precision for rate calculation
    - Tcpdump timestamp accuracy ~ 100us
  - 64-bit registers to store # of lost and out-of-order packets
  - 32-bit register to store running sum of IATs
  - 32-bit register to store running sum of absolute error

- Collected statistics are regularly sent to the control plane
  - Monitoring packets retrieve stored measurements
    - Follows path of generated packets
  - Monitoring packets are tagged with a hardware timestamp for accurate measurements

Concept (III)

- Generated packets contain Ethernet, (VLAN / QinQ), IPv4, UDP, P4TG header

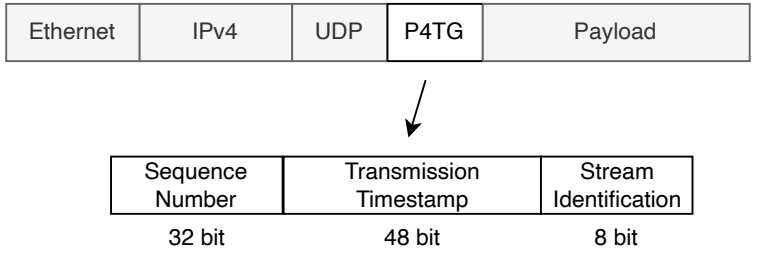

- 32-bit sequence number for packet loss & out-of-order detection
- 48-bit timestamp for RTT calculation
- 8-bit stream identification

Configuration

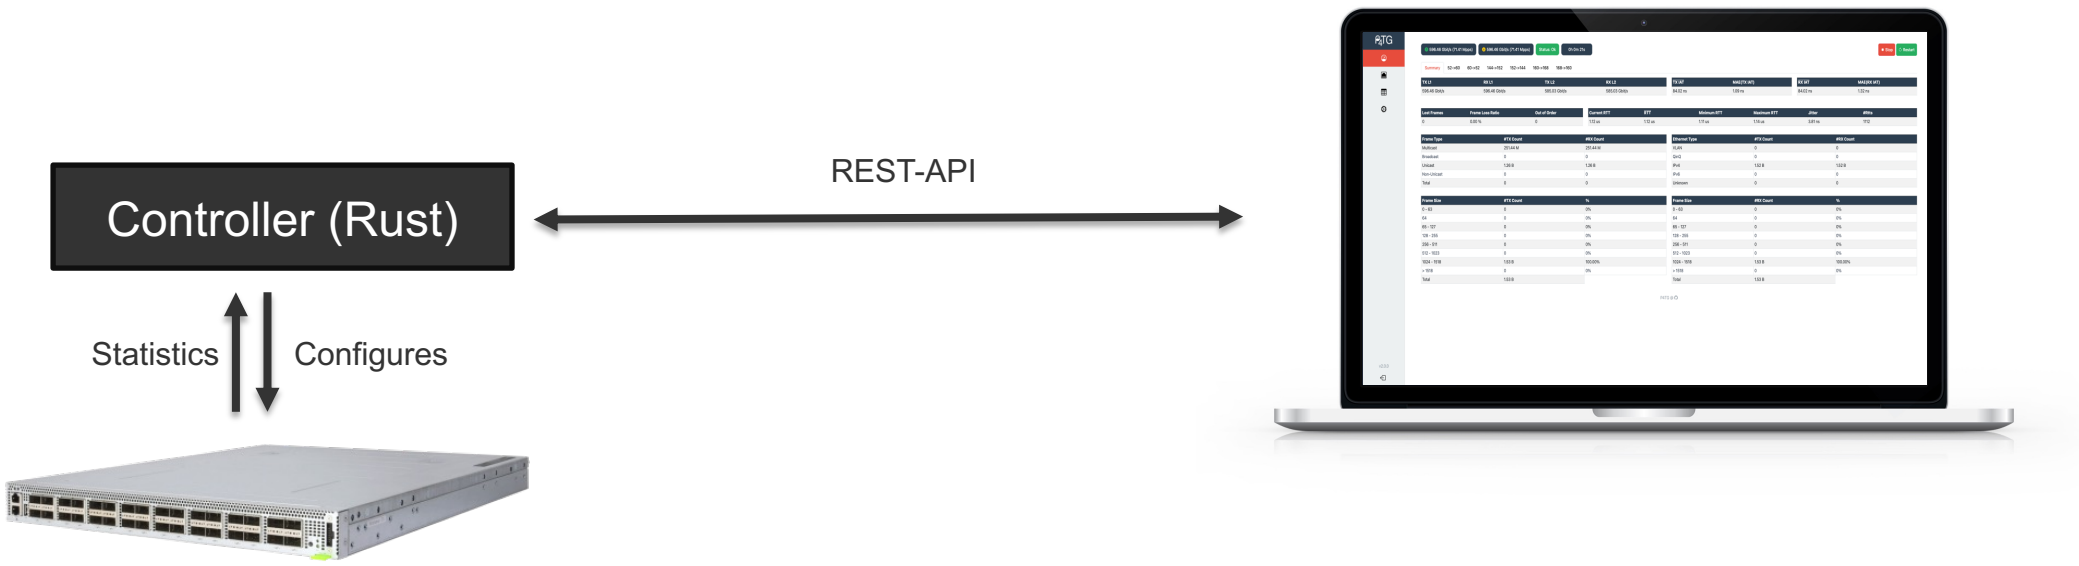

- REST-API may be used to automate tests

Traffic Generation

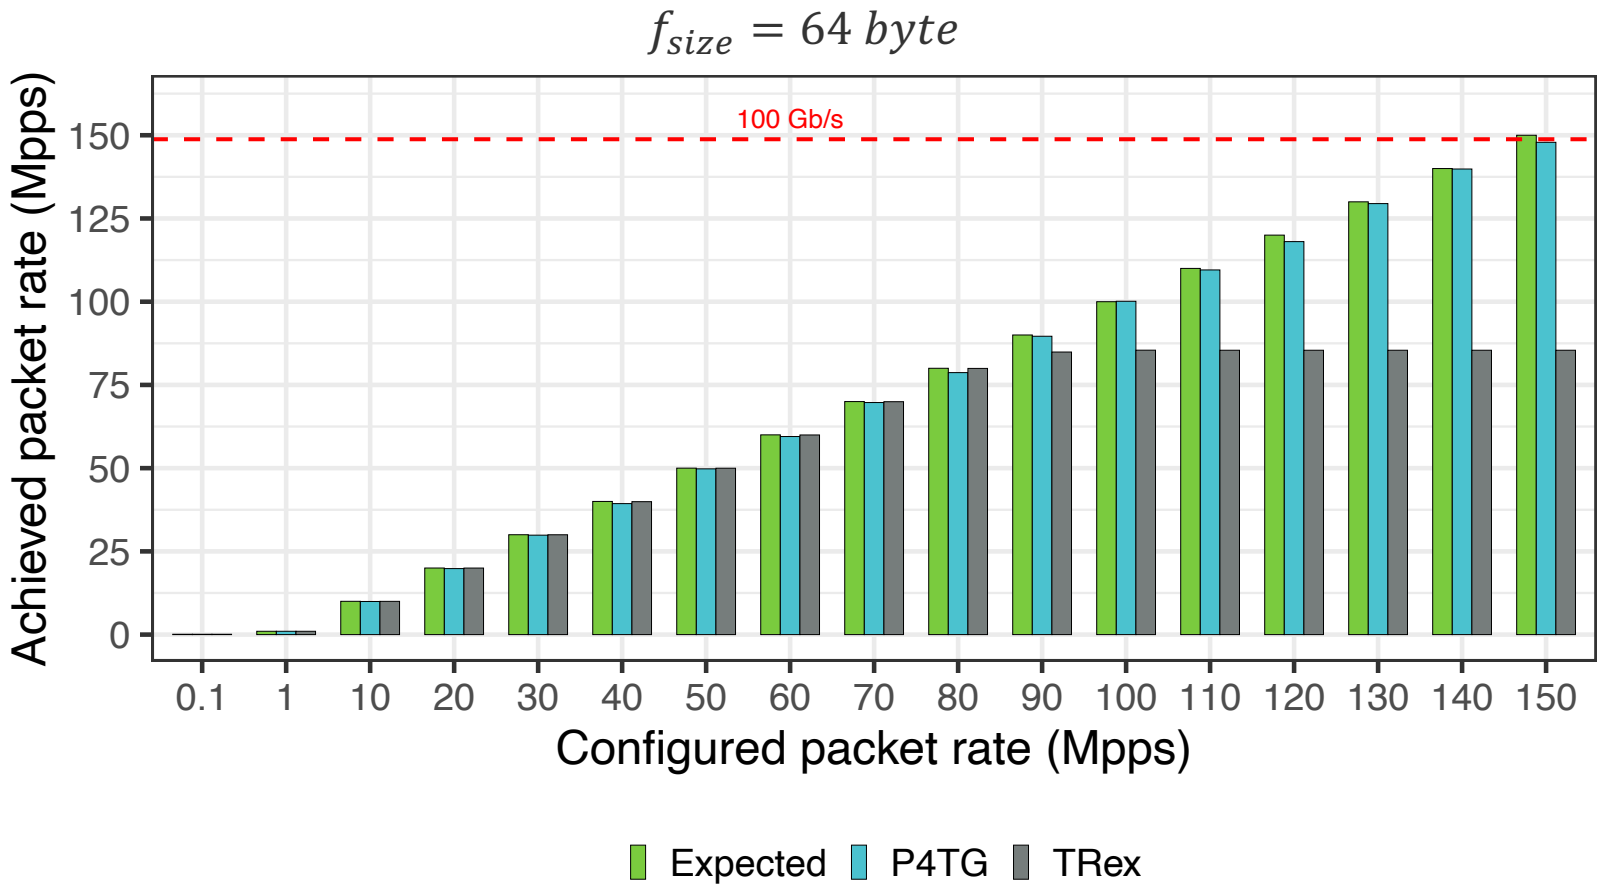

Demo

Conclusion

- P4TG offers traffic generation at high data rates (up to 100 Gbit/s per port)
  - (Possibly) Up to 400 Gbit/s with 2. Gen. Tofino
- Low-cost hardware TG
- Customizable for individual needs
  - Both data and control plane
- https://github.com/uni-tue-kn/P4TG
- https://ieeexplore.ieee.org/document/10048513

Any Questions?

P4TG: 1 Tb/s Traffic Generation for Ethernet/IP Networks

https://github.com/uni-tue-kn/P4TG

Steffen Lindner
University of Tübingen
steffen.lindner@uni-tuebingen.de