

P4-BNG

Central Office Network Functions on Programmable Packet Pipelines

@ Conference on Network and Service Management (CNSM) 2019



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Funded by:




DFG CRC 1053 - MAKI
Multi-Mechanisms Adaptation
for the Future Internet

KuVS Fachgespräch “Network Softwarization”
April 2nd, 2020

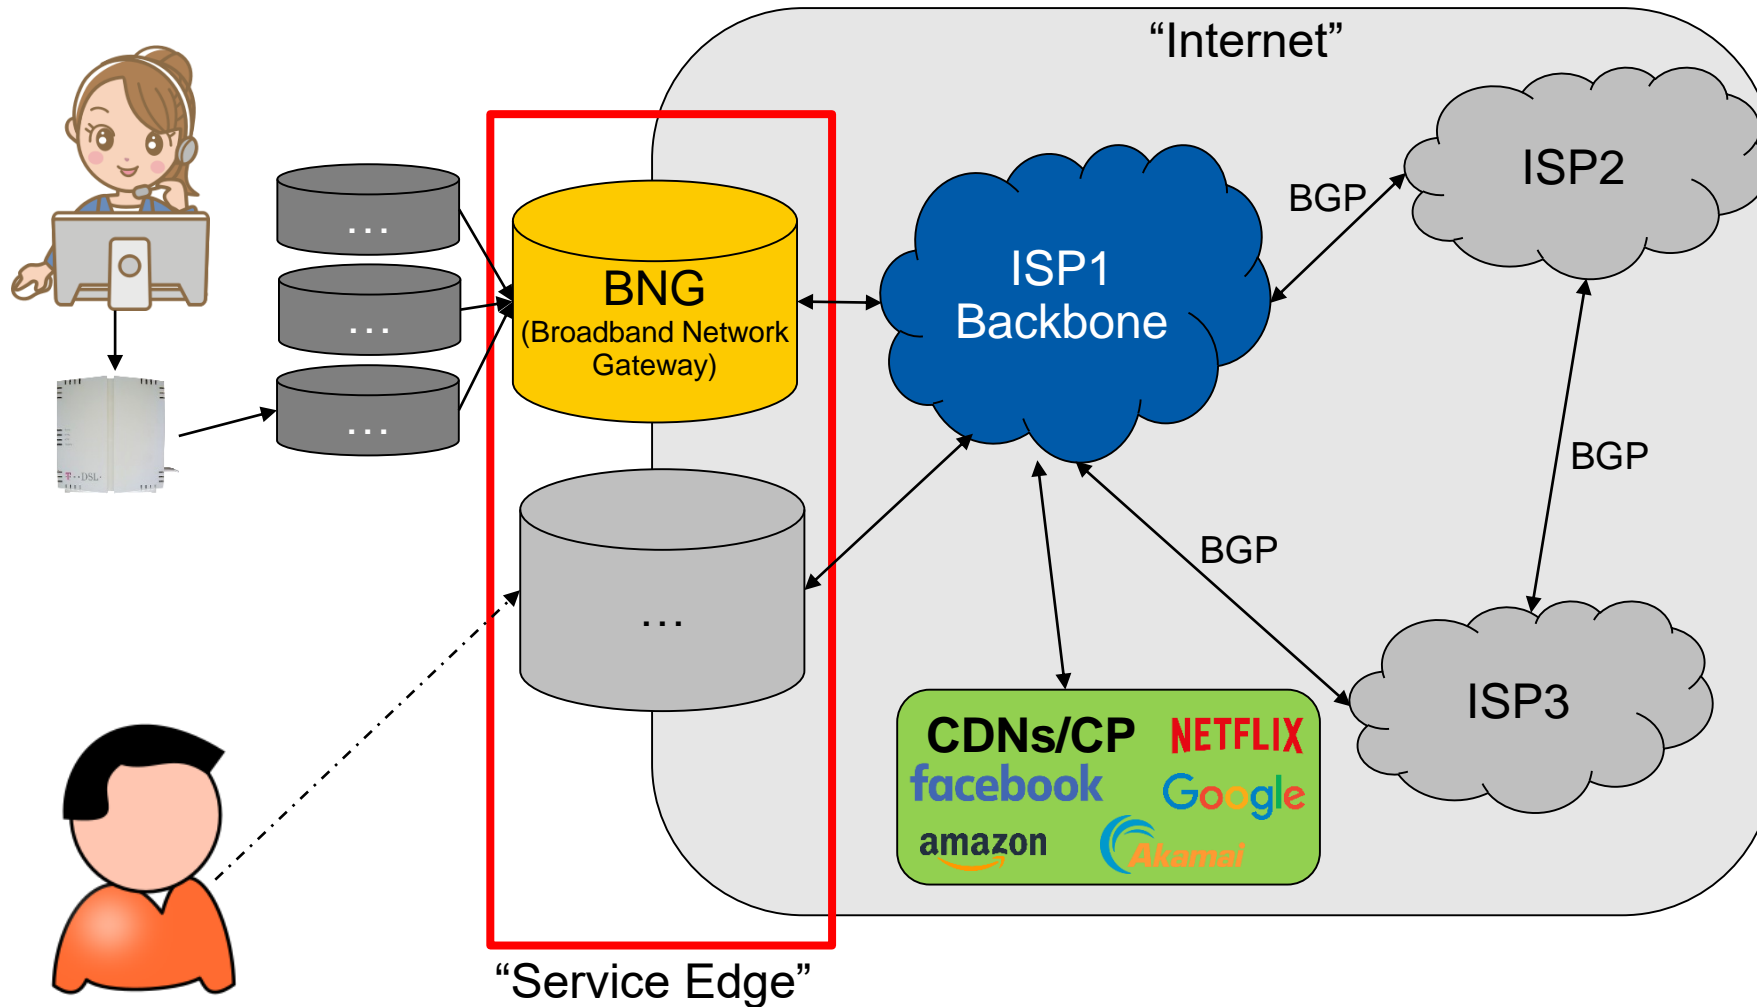
Ralf Kundel

ralf.kundel@kom.tu-darmstadt.de

Leonhard Nobach
Jeremias Blendin
Hans-Joerg Kolbe
Georg Schyguda
Vladimir Gurevich
Boris Koldehofe
Ralf Steinmetz



ISP (access) networks

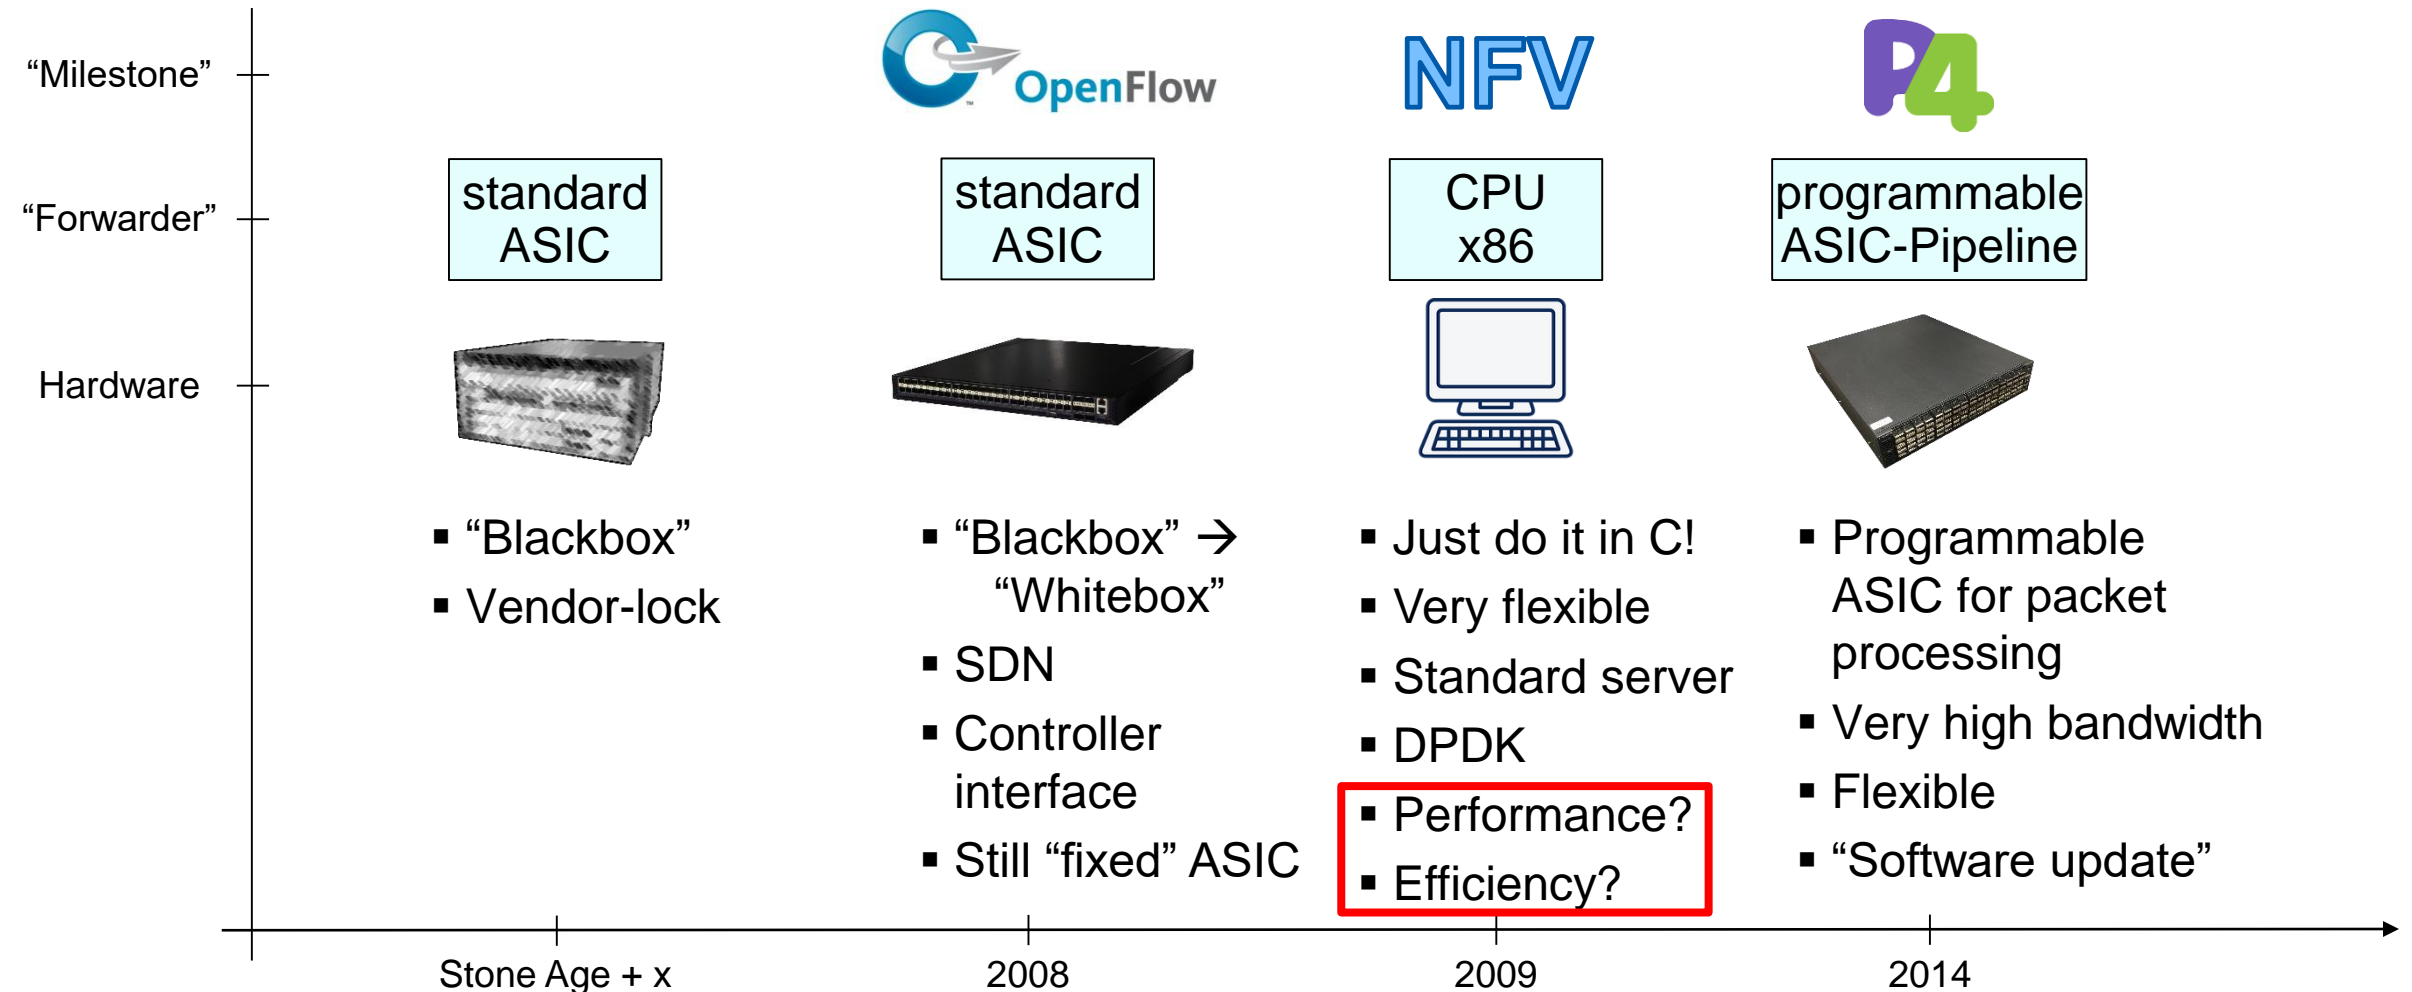


Needed functionality:

- Packet forwarding
- Authentication
- Tunneling
- Accounting
- Multicast
- QoS
- ...


What do Telco's need at the service edge?

History of Data Planes



Why innovation friendly Data Planes?

How to benefit from programmable hardware:



Challenge	“Blackbox” / fixed ASIC	Programmable DP
Time to market	month → years	“software update”
Special requirements	few products	“software update”
New features	buy new hardware	“software update”
Cost	feature dependent	constant low



Can BNG functionality be realized with programmable data plane hardware?



WHAT DOES RELATED WORK SAY?

Open CORD

L.Peterson et al.: "Central Office Re-architected as a Datacenter", 2016



TECHNISCHE
UNIVERSITÄT
DARMSTADT

The CORD Concept


- Mini data centers with **SDN & NFV** at the network edge
- Use commodity data center technology to increase efficiency
- Initial approach: OpenFlow switches and virtual machines on x86
- Operate mobile and residential network access in a single CORD data center
 - R-CORD: Broadband Network Gateway (BNG)
 - M-CORD: Serving Gateway (SGW) + Packet Gateway (PGW) = SPGW
 - E-CORD: Enterprise Customers

Performance?



"economies of scale"

agility of data centers



Bare-Metal Switches And Their Customization ...

L.Nobach et al. @ Local Computer Networks (LCN) 2017



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Bare-metal switch

- “not bundled with an operating system”
- Edge-Core AS5712-54x
- Broadcom Trident2
- OF-DPA



Implementing BNG (BRAS) functionality

- “80% possible”
- E.g. PPPoE missing
- Limited to silicon possibilities
- Very good latency/throughput

standard ASIC
with fixed functionality



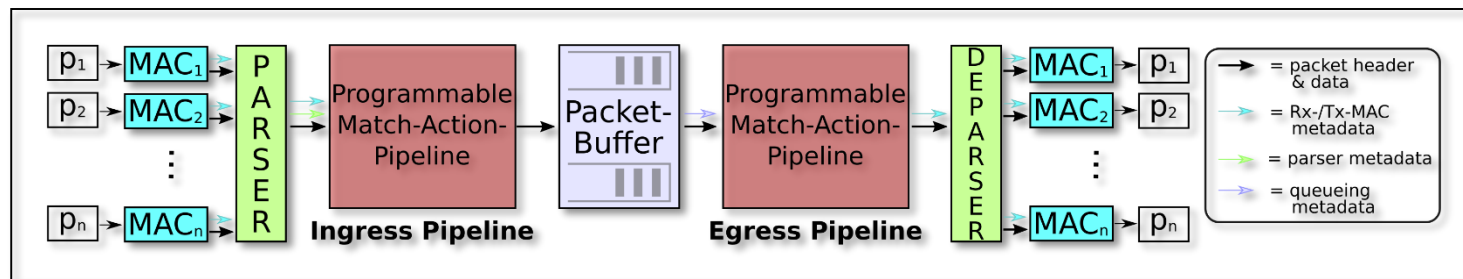
P4

Programming Protocol Independent Packet Processors

Goal: An open source language allowing the specification of packet processing logic

- Protocol independence
- Reconfigure P4 devices by remote “software update”
- No predefined logic!
- “Hardware independent”

P4-BNG?



P4-ASIC



P4-SmartNIC



P4-NetFPGA



What we did:

Analysis and Survey on BNG Requirements

Implement BNG on P4-programmable Data Planes

Evaluation Results

- FPGA
- SmartNIC
- P4-ASIC



Disclaimer:
Vary from country to country!

Functional Requirements BNG

Authentication, Authorization, Accounting

Access control, Security Assurance, DP/CP separation



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Who?

- Residential Gateway (RG) <-> BNG

When?

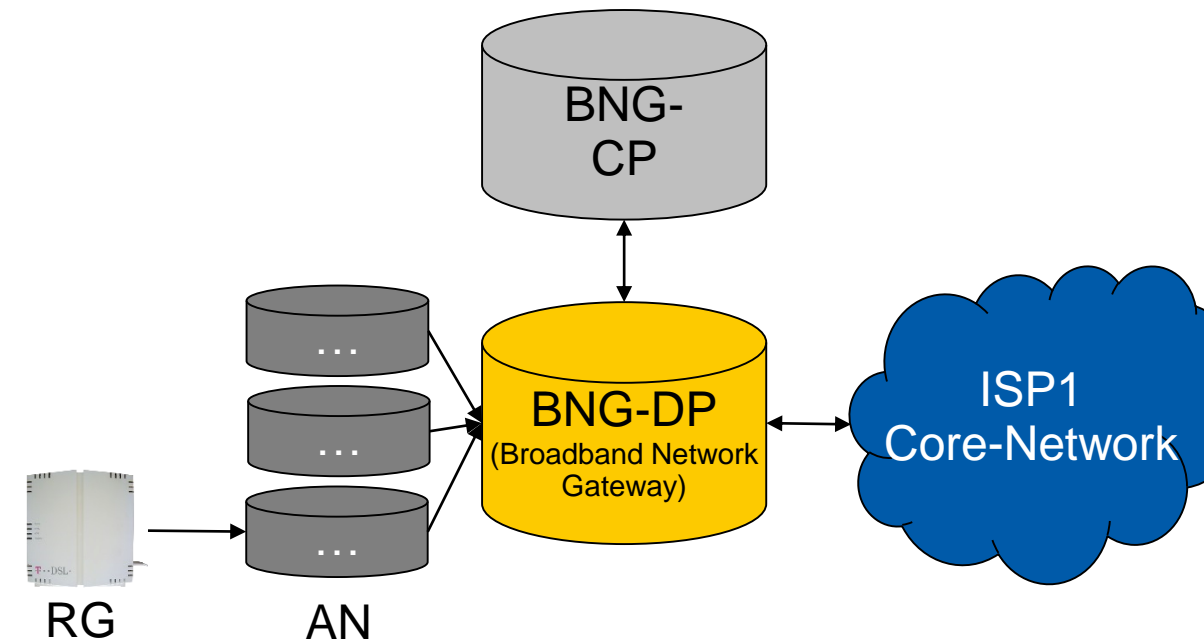
- RG startup
- “blackout”
- Periodically?

How?


- Forward to CP
- Install DP rules
- RPF
- Traffic Separation

Why?

- “Bills payed?”
- Bandwidth Limits?
- Accounting
- “Regulations”
- Customer Separation



Customer Tunneling



Who?

- Residential Gateway (RG) <-> BNG
- “virtual wire” = subscriber line

When?


- Result of AAA

Why?

- Customer Separation
- Prohibit Misuse
- Fairness
- Accounting

How?

- Encapsulation
 - E.g. PPPoE



Traffic rate enforcement



What?

- Downstream traffic
- Upstream traffic


Where?

- At the BNG
 - Downstream shaping
 - Upstream Policing
- RG
 - Upstream shaping?

Why?

- Link speed
- Customer Contract

Bandwidth Limit



Multicast



What?


- Live TV

Where?

- At BNG
- Included in Bandwidth limits

How?

- Useful for live TV only
- MC-groups
 - Customer RG forwards subscriptions to BNG
 - Forward only requested streams
- Offered by ISP

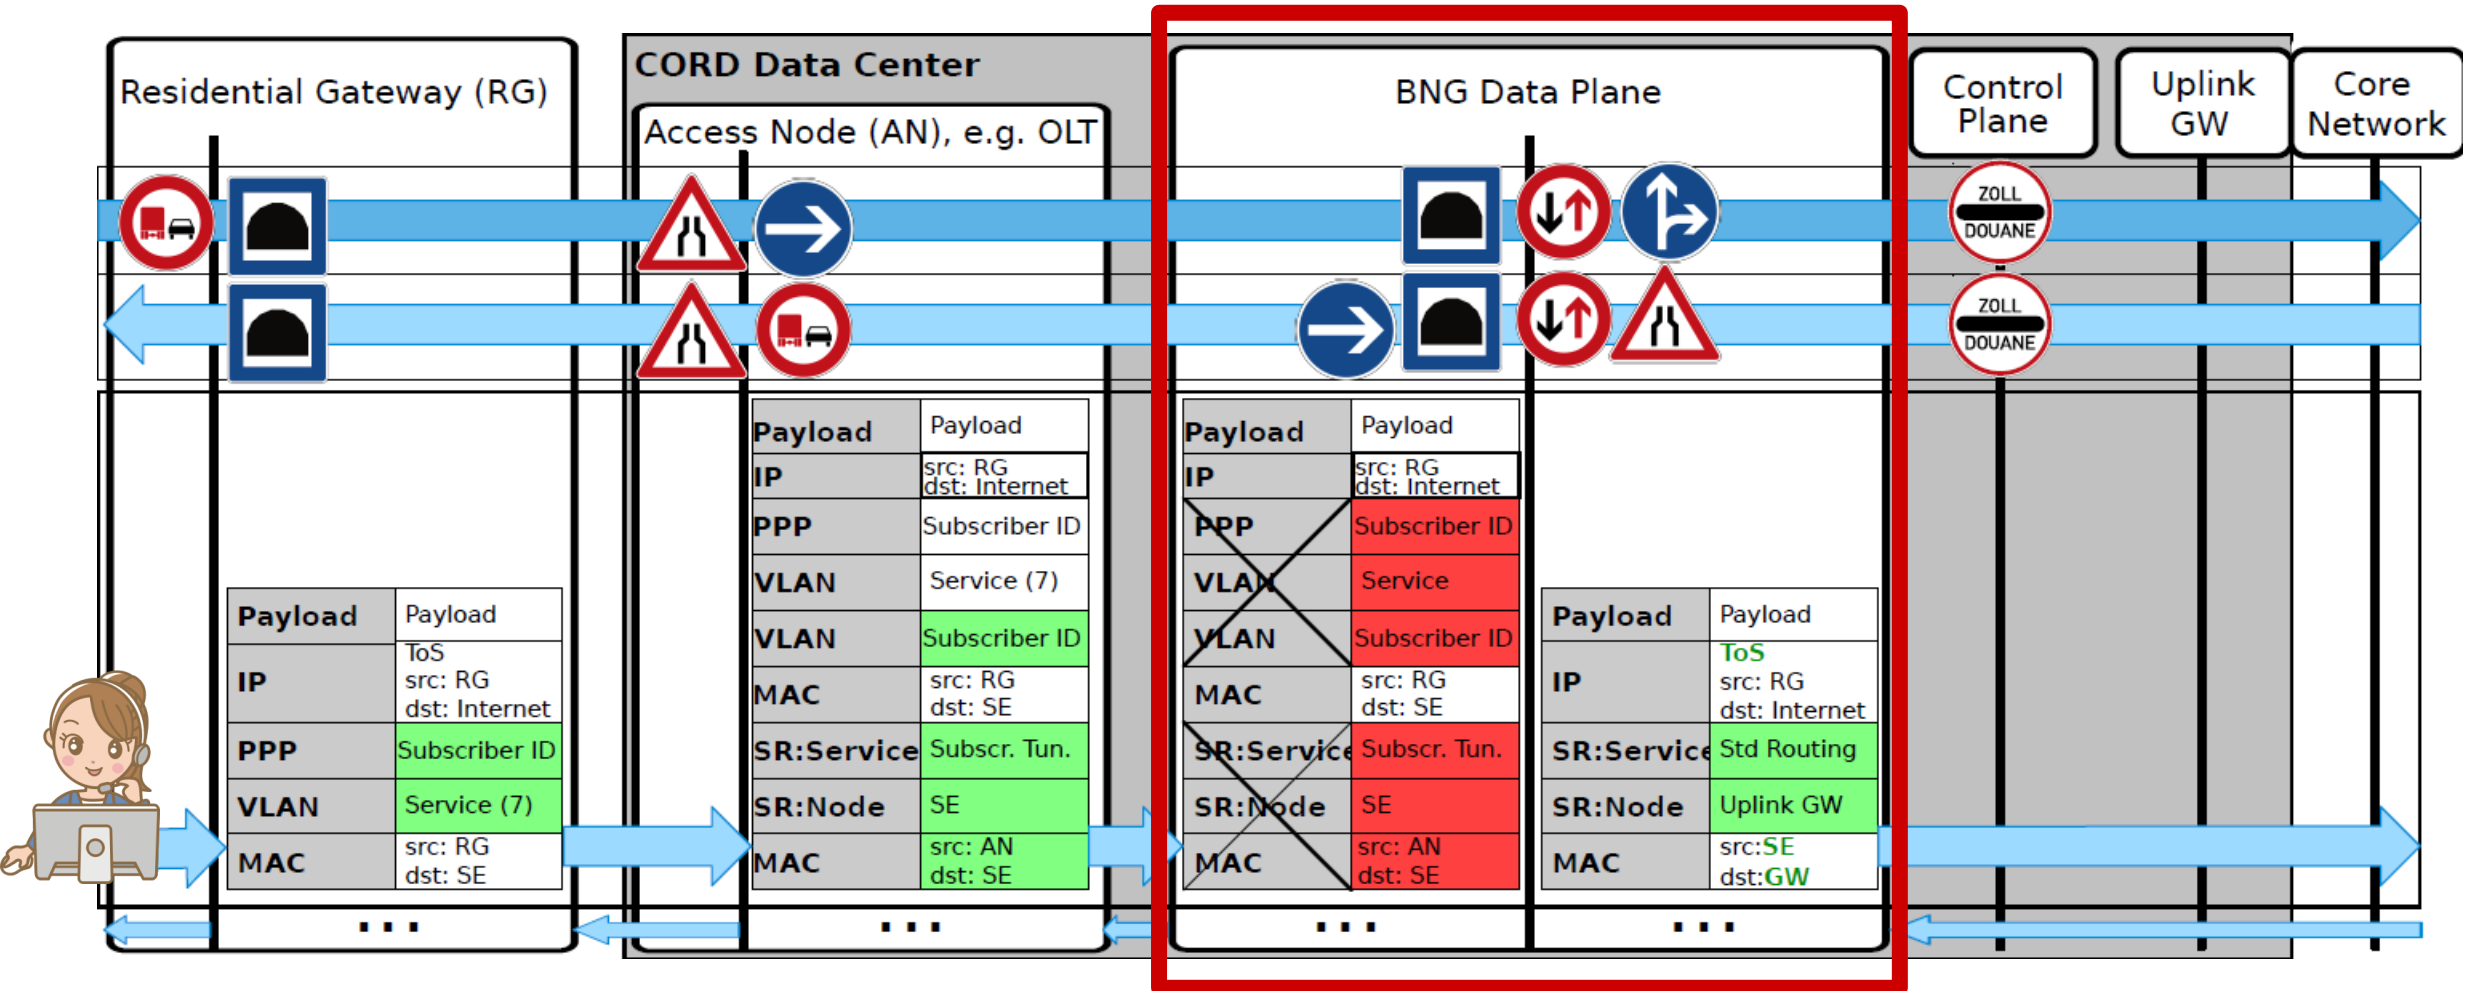




Implementation P4-BNG

P4-BNG header processing – Implemented Design

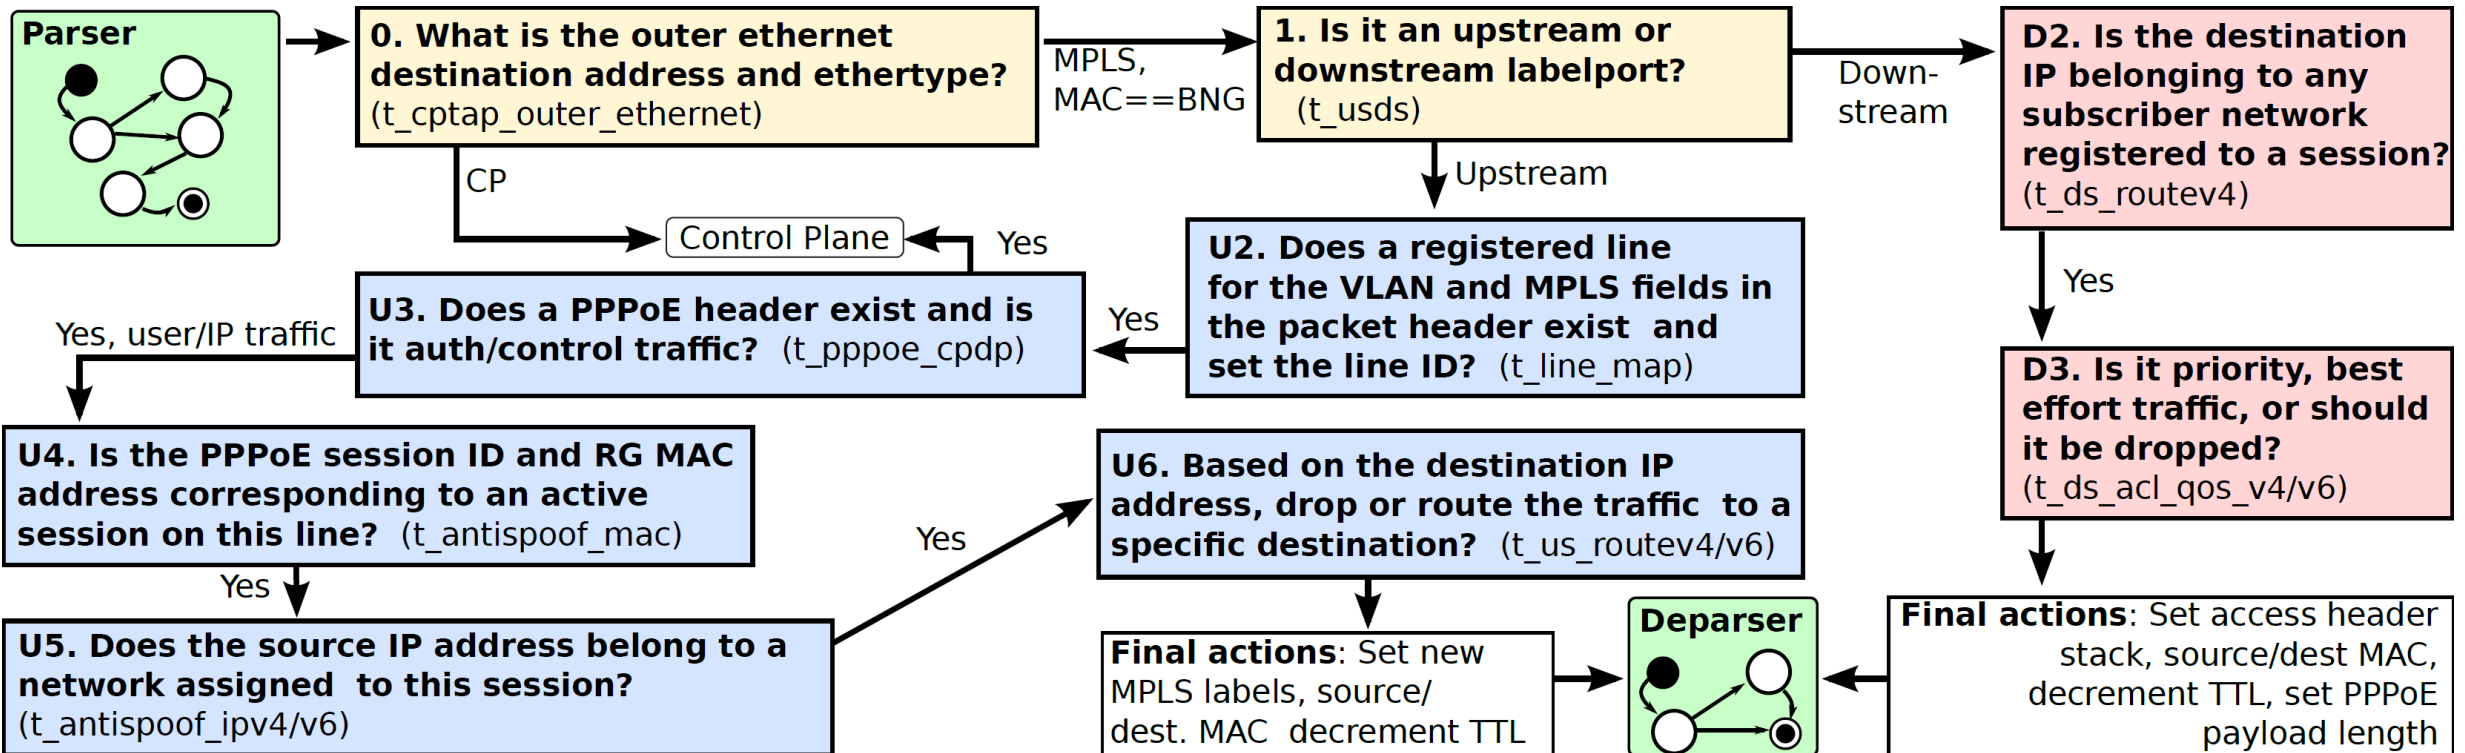
after session setup – “normal case”



P4 control flow



<https://github.com/opencord/p4se>





RESULTS

How to Measure a high performance BNG?

Ingredients

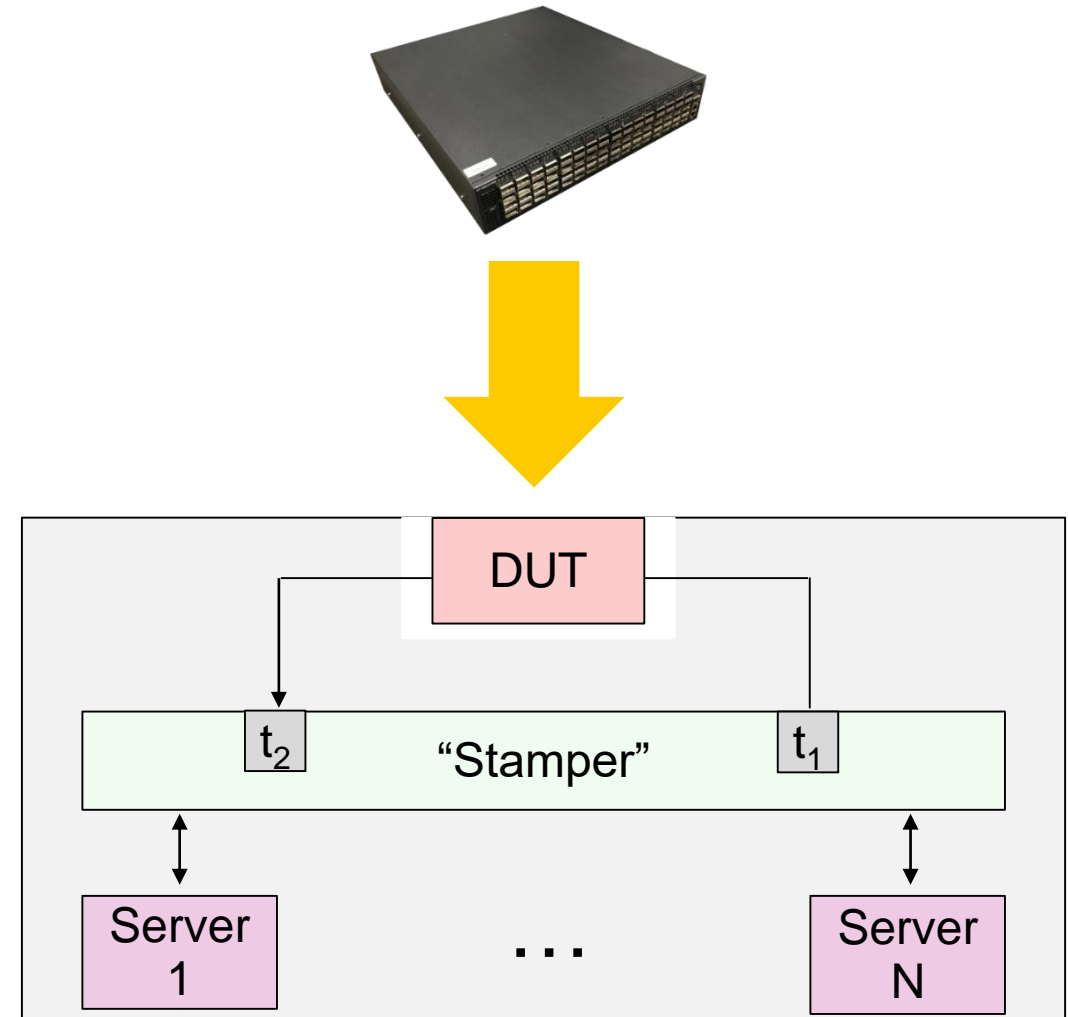
- Device Under Test (P4-BNG)
- P4-Switch as “Stamper”
- N load generation servers

Load Generation

- N servers
- Aggregated by „Stamper“
- Checksum verification can be done in SW

Measurement

- Timestamping in HW
 - t_1 : before DUT
 - t_2 : after DUT
- Loss detection in HW
- Result processing in SW



P4STA: A OpenSource Load Generation Framework

Current State

- OpenSource since 09/2019
- Currently „NDA“ parts missing
- Supported Stamper devices
 - Barefoot Tofino
 - Netronome SmartNICs

Next Steps:

- Further Stamper
 - FPGAs, ...
- ...



<https://github.com/ralfkundel/P4STA>

Configuration Deploy Run Analyze

P4STA configuration environment

KOM – Multimedia Communications Lab

TECHNISCHE UNIVERSITÄT DARMSTADT

[configure switch](#) → [deploy configuration](#) → [run tests](#) → [analyze results](#)

Load/Save Configuration:

[New Configuration](#) netronome_19.02.2021 [Open](#) [Delete](#) [Save](#)

General Settings:

Name of P4 program: <small>Name of the P4 program you want to use.</small>	<input type="text" value="middlebox_v8"/>
Target: <small>Select the stamper device type.</small>	<input type="text" value="delta65p"/>
Stamp Packets: <small>Select the packet type (if supported) to stamp</small>	TCP <input checked="" type="checkbox"/> UDP <input checked="" type="checkbox"/>
SDE <small>Path of SDE directory, e.g. /opt/bf.</small>	<input type="text" value="/opt/bf-sde-8.9.1"/>
Packet Forwarding Mode:	<input type="text" value="Layer 1 (1 server and 1 client only)"/>
Duplication downscale factor: <small>A threshold of 50 causes every 50th packet to be duplicated.</small>	<input type="text" value="20"/>
Packet Generator: <small>Please choose from the available packet generators.</small>	<input type="text" value="iperf3"/>

Configuration of Load Generators:

Loadgen Group	SSH IP (e.g. 172.10.5.9)	SSH Username	Loadgen Iface (e.g. eth0)	Fetch config	Loadgen MAC(e.g. ca:fe:..)	Loadgen IP(e.g. 10.0.0.1)	Set IP	Front Port (e.g. 1/0)	Internal Port	Link Speed (e.g. 10G)	FEC	AN	
Server 1.1	<input type="text" value="172.16.5.12"/>	<input type="text" value="rkundel"/>	<input type="text" value="ens2f0"/> <input checked="" type="checkbox"/>	<input type="text" value="Q"/>	<input type="text" value="3c:fd:fe:9d:0t"/>	<input type="text" value="10.0.10.2/24"/>	<input type="text" value="W"/>	<input type="text" value="20"/>	<input type="text"/>	<input type="text" value="10G"/>	<input type="text"/>	<input type="text"/>	<input type="text" value=""/>
Server 2.1	<input type="text" value="172.16.5.13"/>	<input type="text" value="rkundel"/>	<input type="text" value="ens2f0"/> <input checked="" type="checkbox"/>	<input type="text" value="Q"/>	<input type="text" value="3c:fd:fe:9e:5c"/>	<input type="text" value="10.0.10.3/24"/>	<input type="text" value="W"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="10G"/>	<input type="text"/>	<input type="text"/>	<input type="text" value=""/>

Latency Evaluation



Throughput



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Packet loss in P4

- What's that?
- Except for > 9.99 Gbit/s & “special packet sizes”




Latency Evaluation


Number Subscribers



TECHNISCHE
UNIVERSITÄT
DARMSTADT



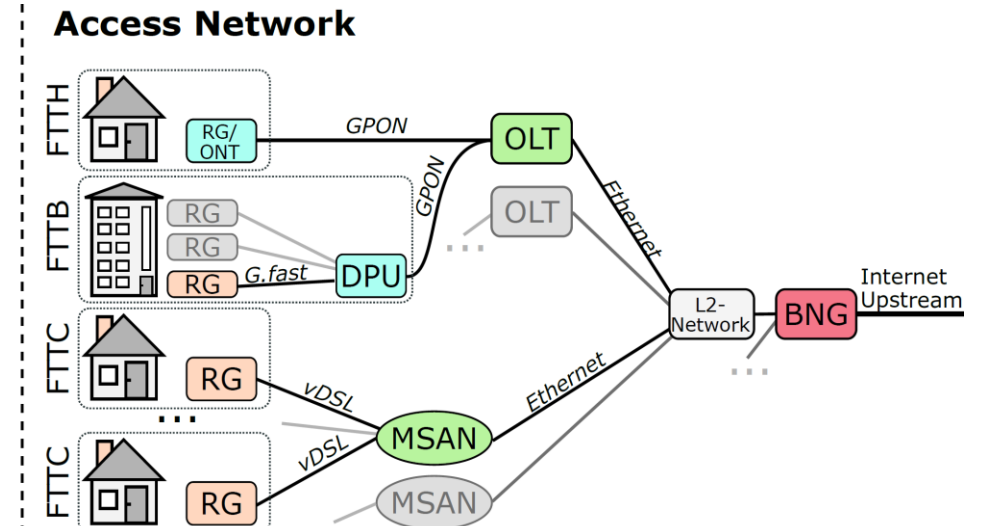

“Hardware either works or not!”



P4-BNG: production ready?

Needed functionality:

- Packet forwarding
- Authentication
- Tunneling
- Accounting
- Multicast
- (H)QoS
- Queues not P4-programmable



Summary and Future Work

Summary

- ISPs have very special requirements
- P4-BNG: It's possible!
- OpenSource available on GitHub



Future Work

- (H)QoS
 - Massive Queueing
 - AQM
 - Hierarchical Scheduling
 -



<https://github.com/opencord/p4se>

Questions & Contact



TECHNISCHE
UNIVERSITÄT
DARMSTADT



Department of Electrical Engineering
and Information Technology
Multimedia Communications Lab - KOM



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Ralf Kundel, M.Sc.

Ralf.Kundel@KOM.tu-darmstadt.de

Rundeturmstr. 10
64283 Darmstadt/Germany
www.kom.tu-darmstadt.de

Phone +49 6151 16-20855
Fax +49 6151 16-29109



<https://github.com/opencord/p4se>