RECENT TRENDS IN TELECOMMUNICATION NETWORKS

SÁNDOR LAKI

DEPT. OF INFORMATION SYSTEMS
FACULTY OF INFORMATICS
ELTE EÖTVÖS LORÁND UNIVERSITY

EMAIL: LAKIS@INF.ELTE.HU
WEB: HTTP://LAKIS.WEB.ELTE.HU

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE
HUNGARY

PROGRAM FINANCED FROM THE NRDI FUND

Paradigm shift in networking - 1

- Closed and proprietary HW and SW
- Implementing 100s of protocols
- Power hungry

Networks are hard to manage
Unreliable, hard to secure
Hard to scale and extend
Slow innovation

- Whitebox switches
  + Servers running the CP
- Merchant silicon
- Linux-based
- Open source

Easier to fix
New features can be added
- Traffic engineering
- Failure handling
- Security

Faster innovation

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Paradigm shift in networking - 2

"This is how I process packets …"

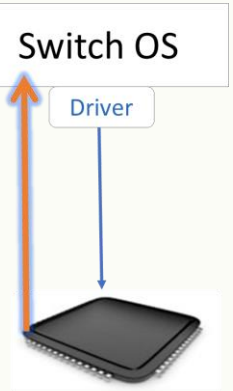

Fixed-function switch

"This is precisely how you must process packets"

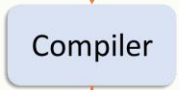

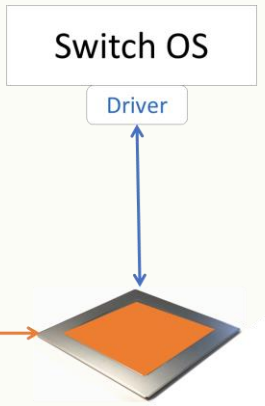

Programmable Switch

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Network as a programmable platform

- Programmable elements (e.g., in P4) at any point of the e2e path

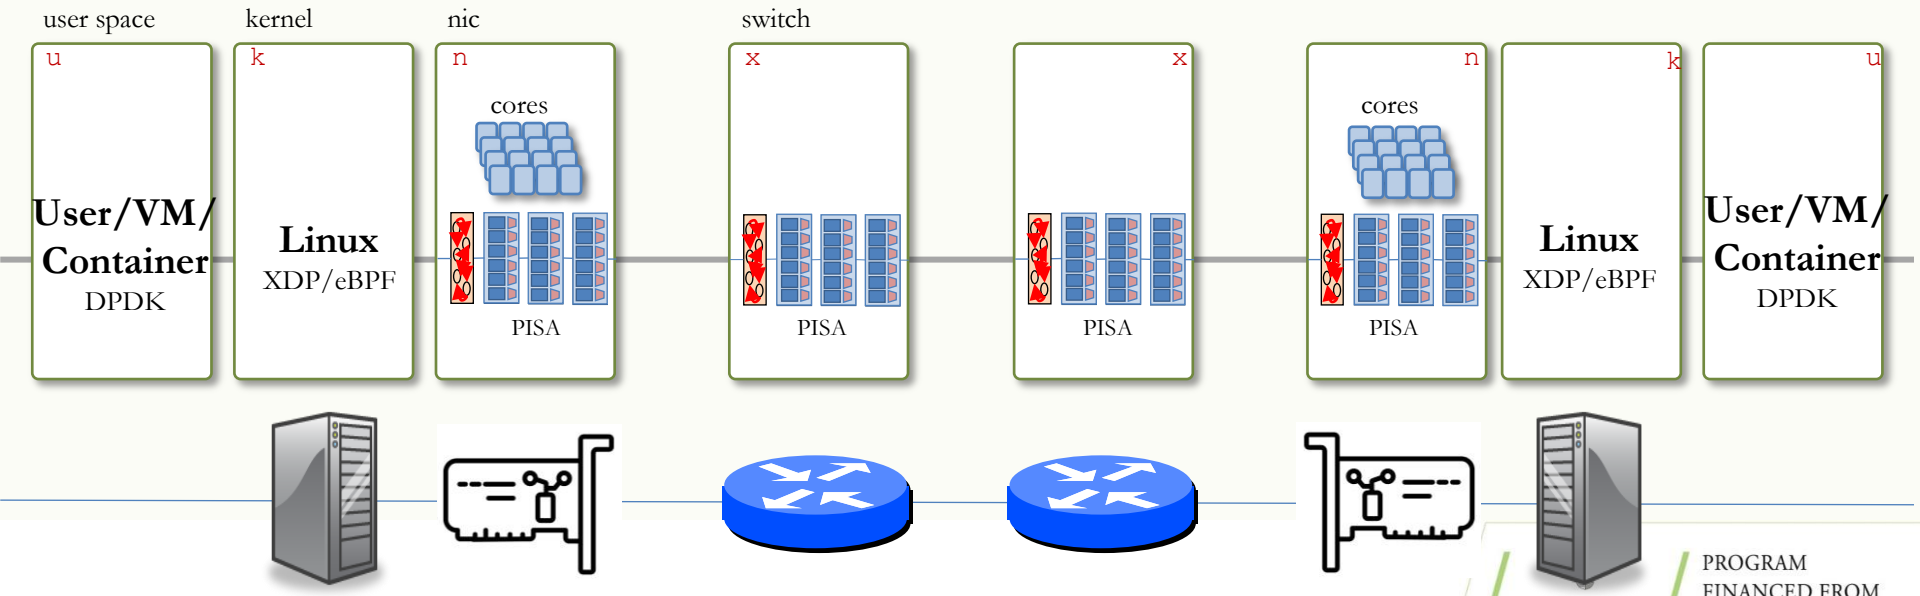

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY

PROGRAM FINANCED FROM THE NRDI FUND

Focus Areas in TKP

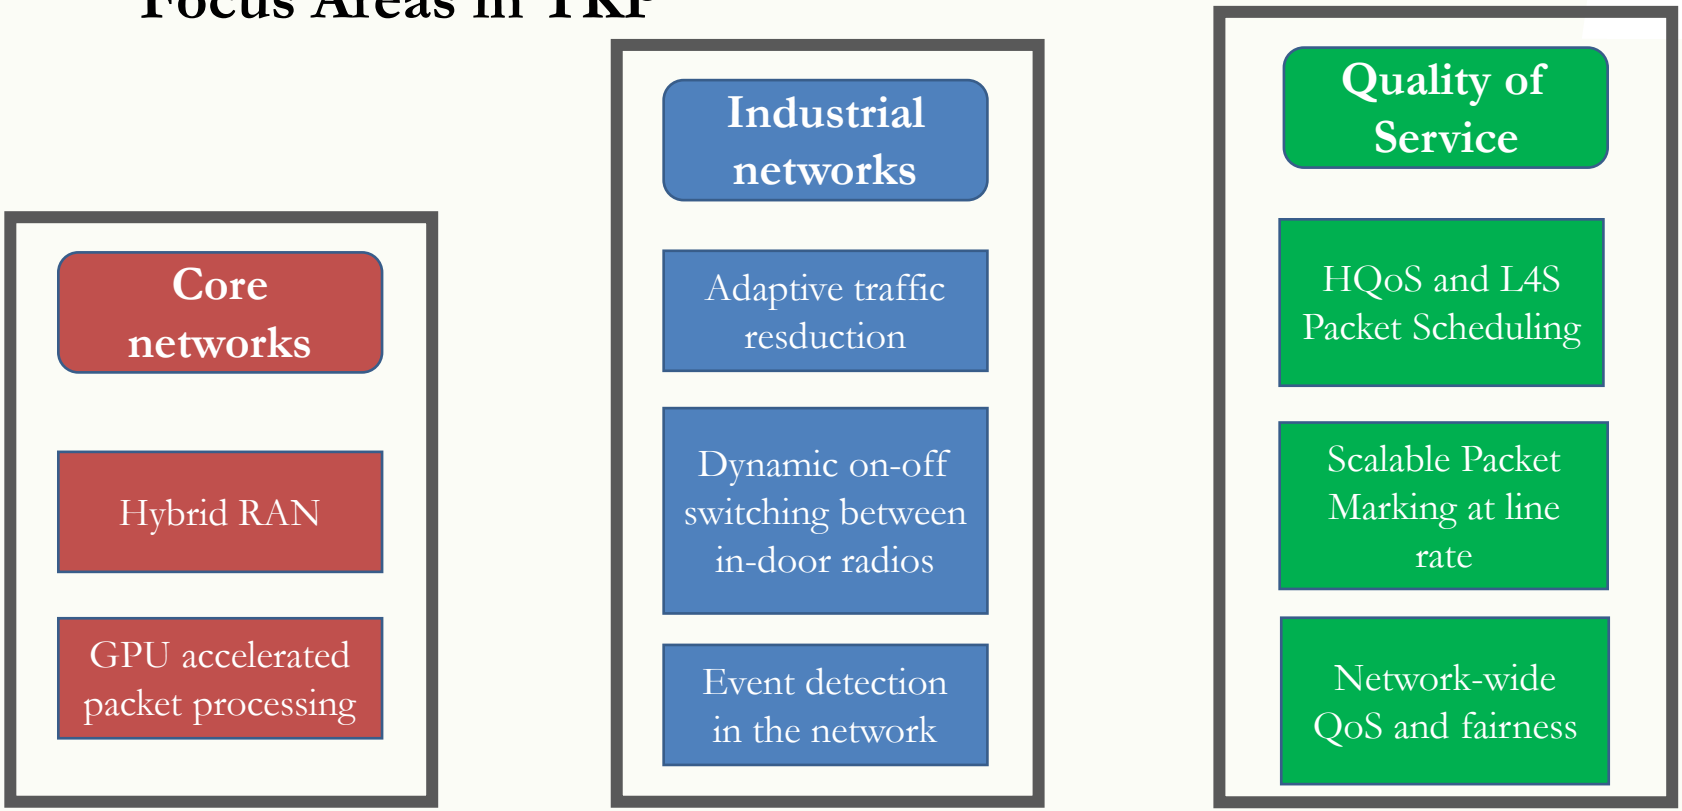

Collaborators in TKP:
Gombos Gergő, Vörös Péter, Szalai-Gindl János, Tejfel Máté, Tóthmérész Lilla, Király Tamás

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Focus Areas in TKP

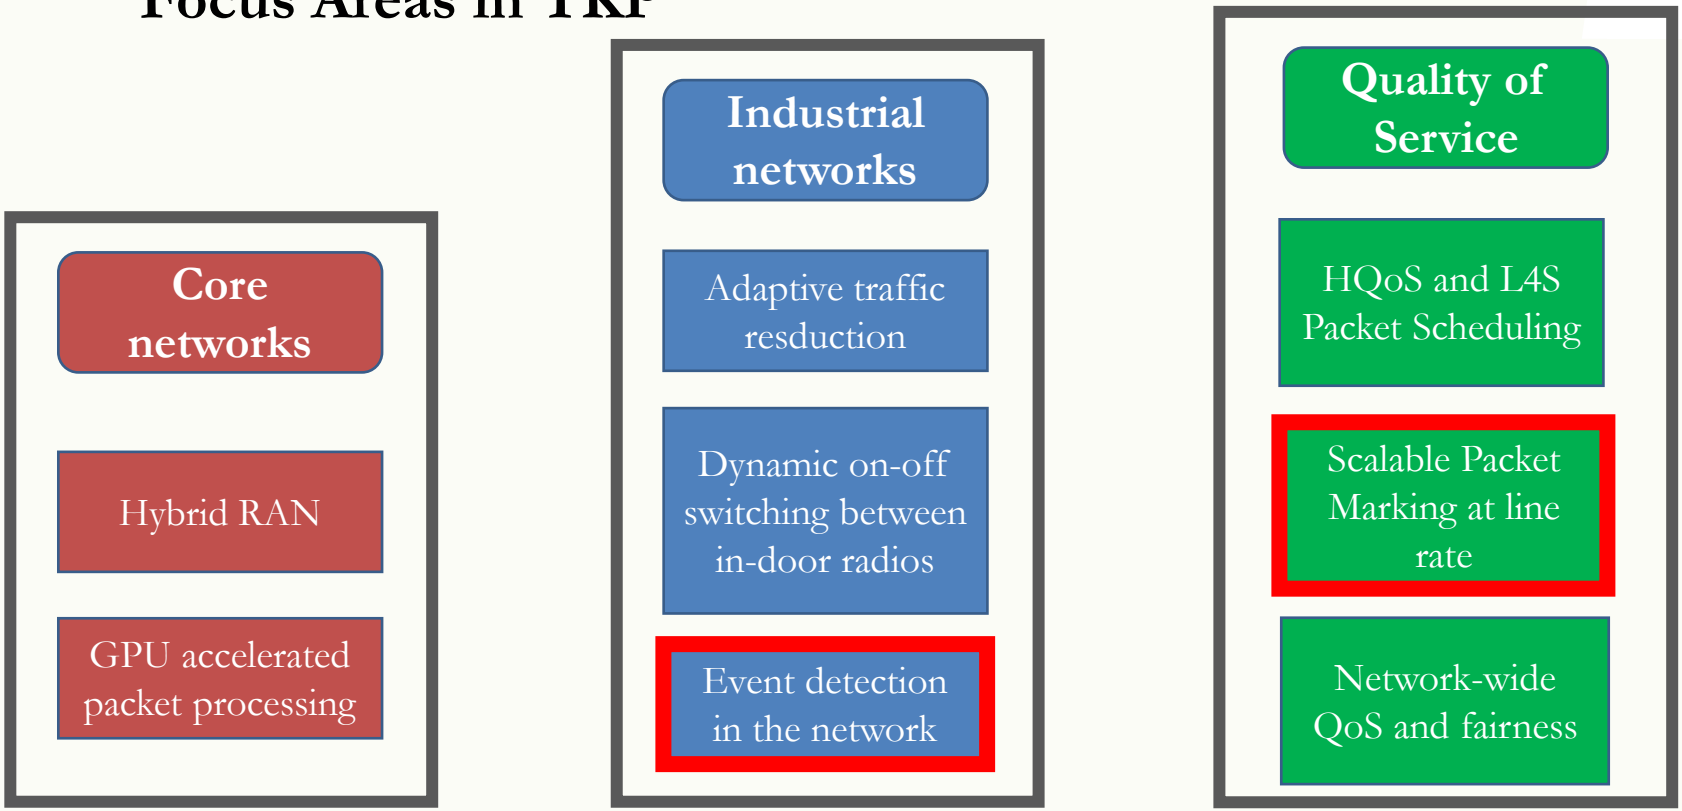

Collaborators in TKP:
Gombos Gergő, Vörös Péter, Szalai-Gindl János, Tejfel Máté, Tóthmérész Lilla, Király Tamás

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Event detection in sensor streams

FastReact: In-network event detection in sensor streams

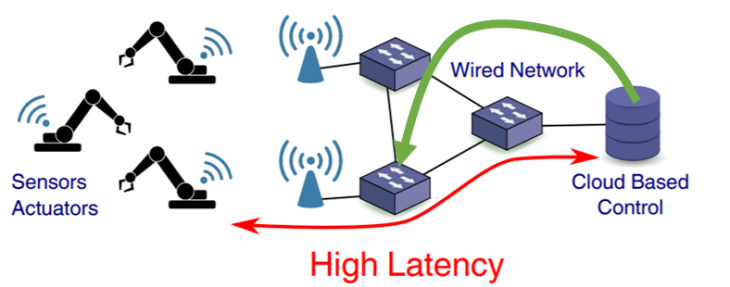

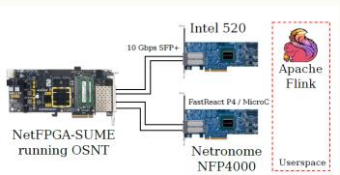

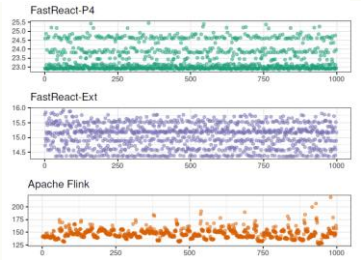

IEEE TNSM – March 2021 (Q1/D1 journal)
Joint work with Andreas Kassler (KAU)

NetReact: Distributed Event Detection with Disaggregated Pipelines

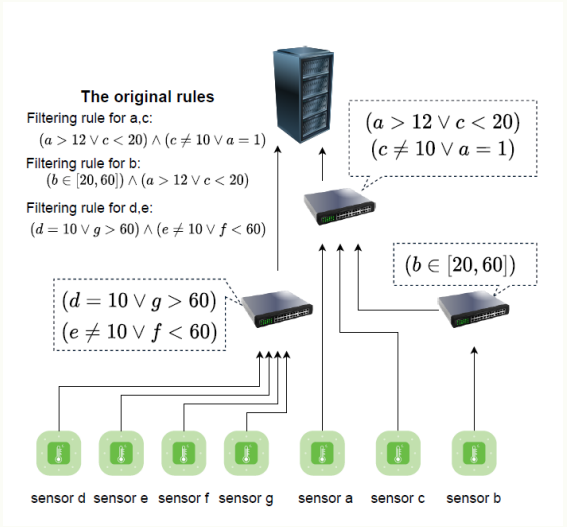

Accepted @ IEEE NetSoft 2022

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

NetReact: Distributed Event Detection with Disaggregated Pipelines

- Run on Intel Tofino HW
  - New algorithm design needed for handling HW constraints
- Sensor values arrive independently
  - Partial evaluation results need to be stored in registers
  - Require mechanisms to avoiding inconsistency
- Algorithm to distribute disjunctions in the network
  - Assuming an optimal topology
  - Rule set represented as a directed graph
  - Strongly connected components moved together

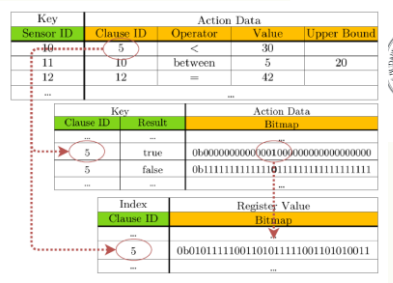

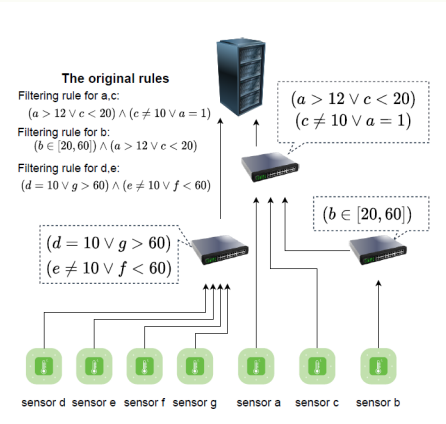

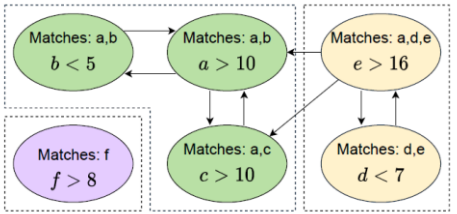

Fig. 5. Strong components and filtering expressions for sensors: $b : (b < 5) \wedge (a > 10)$; $a, c : (a > 10) \wedge (e > 16) \wedge (c > 10)$; $d, e : (d < 7) \wedge (e > 16)$

NetReact: Distributed Event Detection with Disaggregated Pipelines

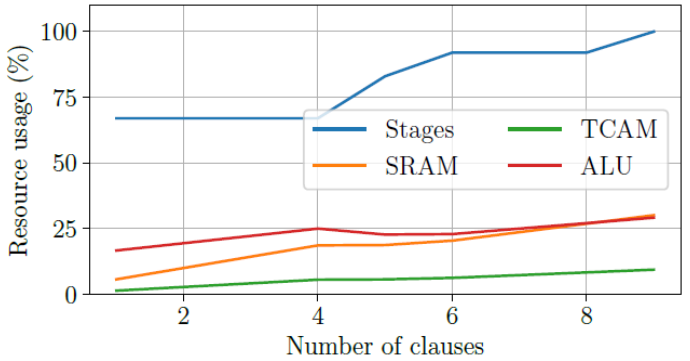

- Computation hungry, but has low memory fingerprint
- Scales linearly with the number of switches

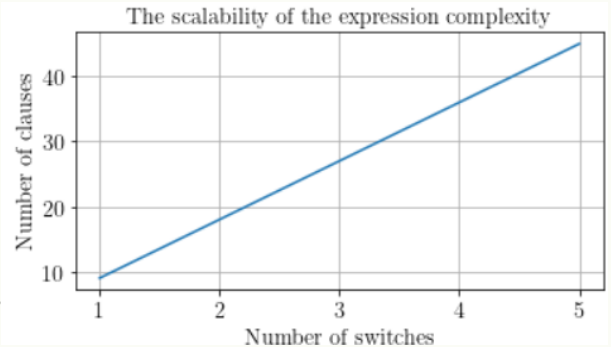

- More details in our paper:

Cs. Györgyi, K. Kecskeméti, H. Mallouhi, P. Vörös, S. Laki: NetReact: Distributed Event Detection in Sensor Data Streams with Disaggregated Packet Processing Pipelines, IEEE International Conference on Network Softwarization (IEEE NetSoft'22), 27 June–1 July 2022, Milan, Italy

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Packet Tagging-based QoS framework

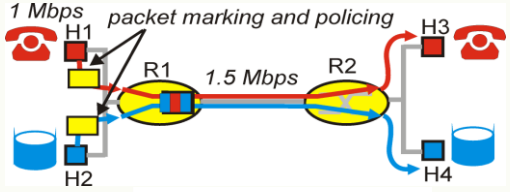

Core-stateless Packet Marker in Intel Tofino

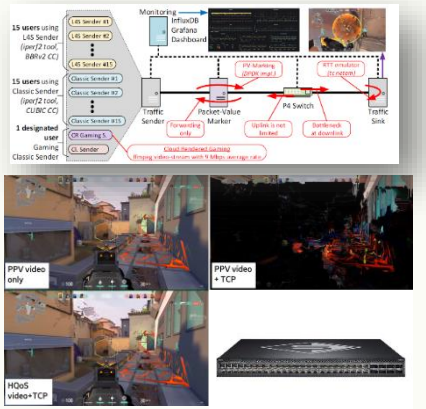

IEEE INFOCOM'21 – BEST DEMO AWARD
IEEE TNSM (Q1/D1 journal) 2022

Quality of Service

Scalable Packet Marking in HW switches

HQoS and L4S Scheduling In HW switches

National Research, Development and Innovation Office Hungary

PROGRAM FINANCED FROM THE NRDI FUND

Our approach

- **Packet Marker** at the edge of the network
  - **Stateful, but highly *distributed***
  - Two markings:
    - Packet Value: expressing the importance of packet
    - Delay Class: expressing a strict delay requirement
- **Resource Nodes** (e.g. routers) drop/schedule packets according to packet markings
  - **Stateless and simple**
  - Solely use the markings:
    - Packet Value: decision on which packet to drop if needed
    - Delay/Traffic class: how to schedule the packet

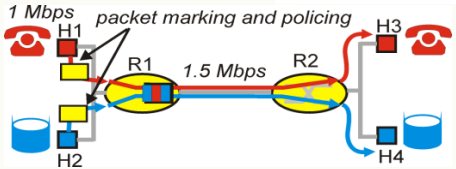

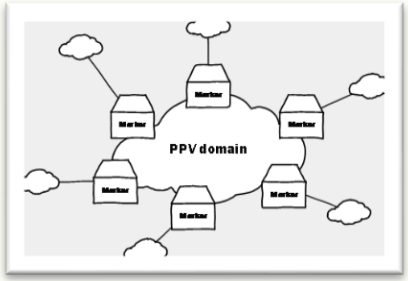

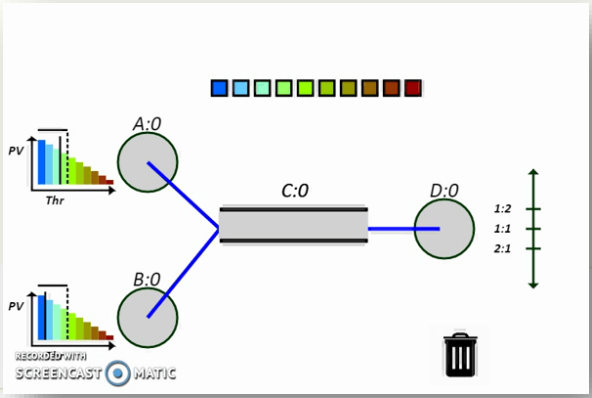

Packet Marking

- Independent markers for each traffic aggregate (TA)

- **Packet Value** based on a given *TVF*
  - Continously measures the rate *R* of TA
  - $R_{current} = random(0,R)$
  - $V = TVF(R_{current})$

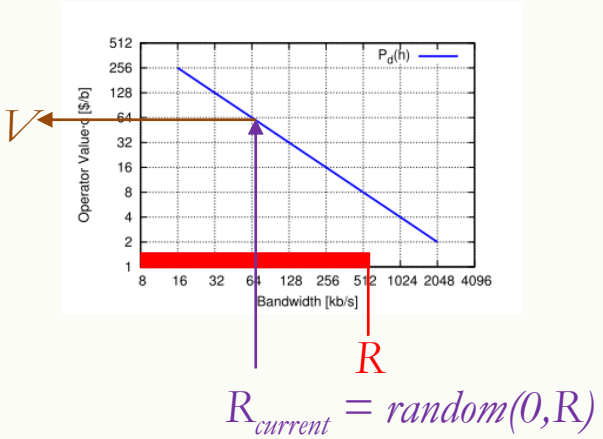

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Scalable Packet Marking in HW switches

- Problems in HW data planes
  - Rate measurements – LPF
  - Limited random number generator
  - No floating point arithmetic
  - Long computational chains not possible
- Exponential binning and logarithmization trick
  - Assuming TVF has an exponential decay
  - Solved in prefilled match-action tables

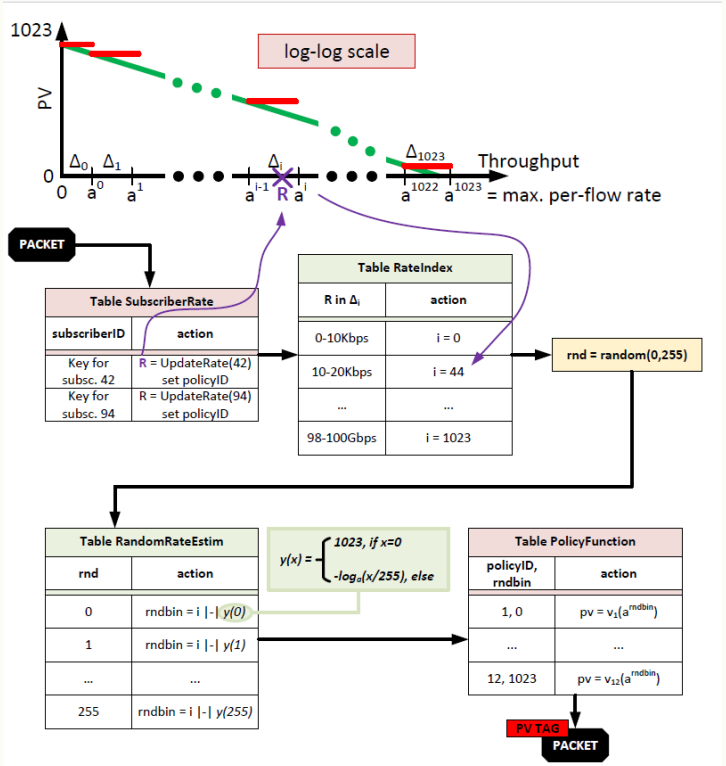

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Scalable Packet Marking in HW switches

- Our non-optimized prototype
  - Up to 35K subscribers
  - 100Gbps bottleneck >200Gbps load
  - Almost perfect weighted fairness

8000 Silver + 2000 Gold subscribers

- Testbed

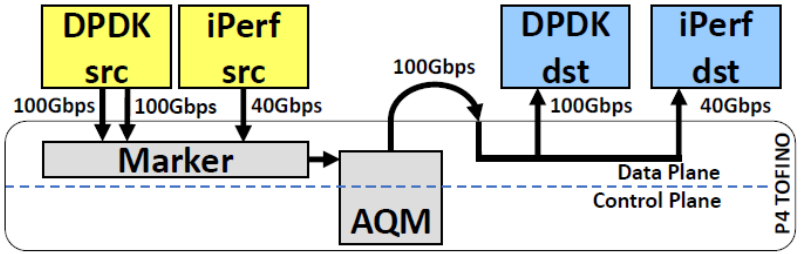

Demo submitted to ACM SIGCOMM 2022

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY
PROGRAM FINANCED FROM THE NRDI FUND

Scalable Packet Marking in HW switches

- Our non-optimized prototype
  - Up to 35K subscribers
  - 100Gbps bottleneck >200Gbps load
  - Almost perfect weighted fairness

28000 Silver + 7000 Gold subscribers

- Testbed

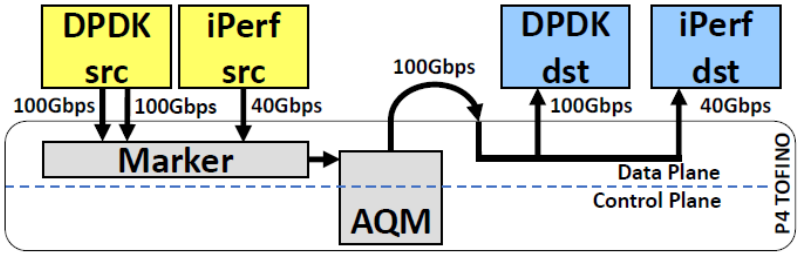

Demo submitted to ACM SIGCOMM 2022

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

2019-2022 **Publications**

- **Accepted journal papers:**
  - **Q1: 4 [D1: 3], Q2: 1**
  - ACM/IEEE TNET, IEEE TNSM, SIGCOMM CCR, IEEE ACCESS
- **Journal papers under review:**
  - **Q1: 5 [D1: 2], Q2: 1**
  - ACM/IEEE TNET, IEEE ACCESS, IEEE TMC, SIGCOMM CCR
- **Accepted conference papers: 18**
- **Conference papers under submission: 3**

- Selected papers:
  - High impact conferences:
    - G. Gombos, M. Mouw, S. Laki, C. Papagianni, K. De Schepper: Active Queue Management on the Tofino programmable switch: The (Dual)PI2 case, IEEE International Conference on Communications 2022 (IEEE ICC'22) Next-Generation Networking and Internet Symposium, 16-20 May 2022, Seoul, South Korea.
    - S. Laki, Cs. Györgyi, J. Pető, P. Vörös, G. Szabó: In-Network Velocity Control of Industrial Robot Arms, 19th USENIX Symposium on Networked Systems Design and Implementation (NSDI'22), April 2022, Renton, WA, USA. https://www.usenix.org/conference/nsdi22/presentation/laki
  - Best Demo Award @ RankA*:
    - F. Fejes, Sz. Nádas, G. Gombos, S. Laki: A Core-Stateless L4S Scheduler for P4-enabled hardware switches with emulated HQoS, IEEE International Conference on Computer Communications (IEEE InfoCom'21), 10-13 May 2021, Virtual Conference. https://ieeexplore.ieee.org/document/9484581
  - Q1/D1 journals:
    - F. Fejes, Sz. Nádas, G. Gombos, S. Laki: DeepQoS: Core-Stateless Hierarchical QoS in Programmable Switches, IEEE Transactions on Network and Service Management (IEEE TNSM) journal (Q1, IF: 5.93), Volume: tba, Issue: tba, Page(s): 1-20, 2022, doi: 10.1109/TNSM.2022.3152017
    - S. Laki et al.: Core-Stateless Forwarding with QoS Revisited: Decoupling Delay and Bandwidth Requirements, IEEE/ACM Transactions on Networking (IEEE/ACM ToN) journal (Q1, IF: 3.315), vol. 29, no. 2, pp. 503-516, April 2021, doi: 10.1109/TNET.2020.3041235

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY
PROGRAM FINANCED FROM THE NRDI FUND

2019-2022 **Grants and project proposals**

- MTA Bolyai scholarship
  - Closed in 2021 Sept with outstanding qualification
- OTKA
  - 1 accepted proposal
  - 2 rejected proposals
- EU COST Action
  - 2 submissions as Main proposer
    - **Both rejected**

- EU H2020 5GinFIRE OC-2
  - P4in5G: 75K EUR for 6M
- EU H2020 NGI Pointer OC-2
  - P4EDGE: 200K EUR for 12M
- EU HE SNS-6G Stream-B
  - Submitted, under eval.
  - 6M EUR budget 14 partners

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Other activities and results

- **TechTalks at P4 Workshop 2021 and 2022**

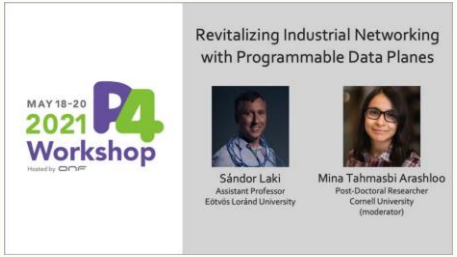

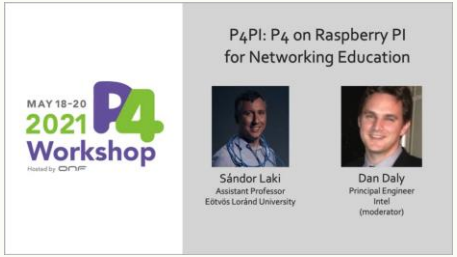

- **Talk at the QoS WS'21 of Internet Architecture Board**
  - Joint with Telefonica and Ericsson
- **Project with ONF P4 Education WG**
  - P4PI: P4 on Raspberry Pi for Networking Education
    - Noa Zilberman (Oxford), Robert Soulé (Yale), Fernando Ramos (Uni. Lisbon)
  - **Hackathon at ACM SIGCOMM 2021** (Rank A*)
  - **Accepted hackathon at ACM SIGCOMM 2022**

- Invitation to a **Dagstuhl Seminar**
  - Towards More Flexible and Automated Communication Networks, Nov. 2022

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY
PROGRAM FINANCED FROM THE NRDI FUND

Others

- Cooperation with the mathematical optimization RG
  - Lilla Tothmérész, Tamás Király
- External collaborators with joint papers
  - Andreas Kassler,Karlstad University (Sweden)
  - Christian E. Rothenberg, University of Campinas (BR)
  - Chrysa Papagianni, University of Amsterdam (NL)
  - Noa Zilberman, University of Oxford (UK),
  - Robert Soulé, Yale University (USA),
  - Koen de Schepper, Nokia Bell Labs Anwerp (BE)
  - Stefan Schmid, TU Berlin (DE)
  - Luis M. Contreras, Telefonica (S)
  - Szilveszter Nádas, Gergely Pongrácz, Géza Szabó, Ericsson Research (HU)
- Ongoing discussions on
  - HQoS with UC Berkeley (USA)
  - in-network computing with Pensando (AMD)
- TPC memberships:
  - USENIX ATC'22 (Rank A), IEEE CCNC 2020-2022 (Rank B), IEEE NetSoft 2021 (Rank B), IEEE CloudNet 2021-2022, EUROP4 WS 2020-2021

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Special thanks to

My closer colleagues: Gergő Gombos, Péter Vörös, János Szalai-Gindl, Dhulfiqar A. AlWahab

My PhD students: Dániel Varga, Csaba Györgyi, Ferenc Fejes, Dávid Kis, Hiba Mallouhi, Ákos Rudas

Q&A

WEB: HTTP://LAKIS.WEB.ELTE.HU

Az Alkalmazásiterület-specifikus nagy megbízhatóságú informatikai megoldások című projekt a Nemzeti Kutatási Fejlesztési és Innovációs Alapból biztosított támogatással, a Tématerületi kiválósági program (TKP2020-NKA-06, Nemzeti Kihívások Alprogram) finanszírozásában valósult meg.

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY

PROGRAM FINANCED FROM THE NRDI FUND