DEEPLY PROGRAMMABLE NETWORKS

SEPTEMBER 2020 – JUNE 2021

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 …"

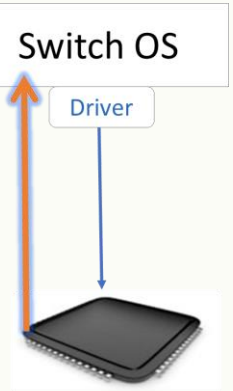

Fixed-function switch

"This is precisely how you must process packets"

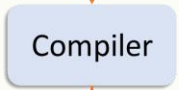

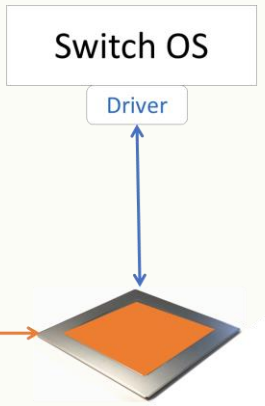

Programmable Switch

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

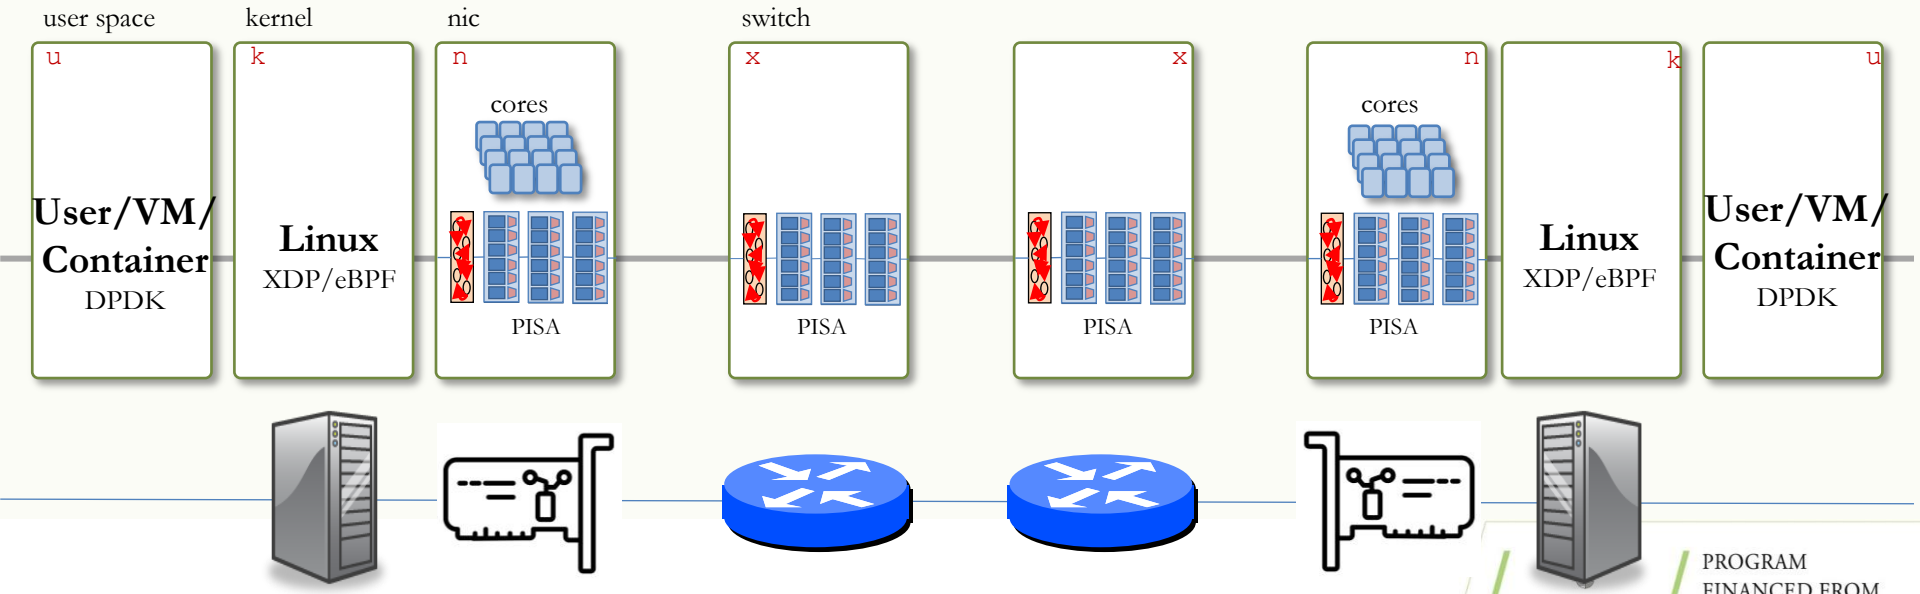

PROGRAM
FINANCED FROM
THE NRDI FUND

September 2020 – June 2021

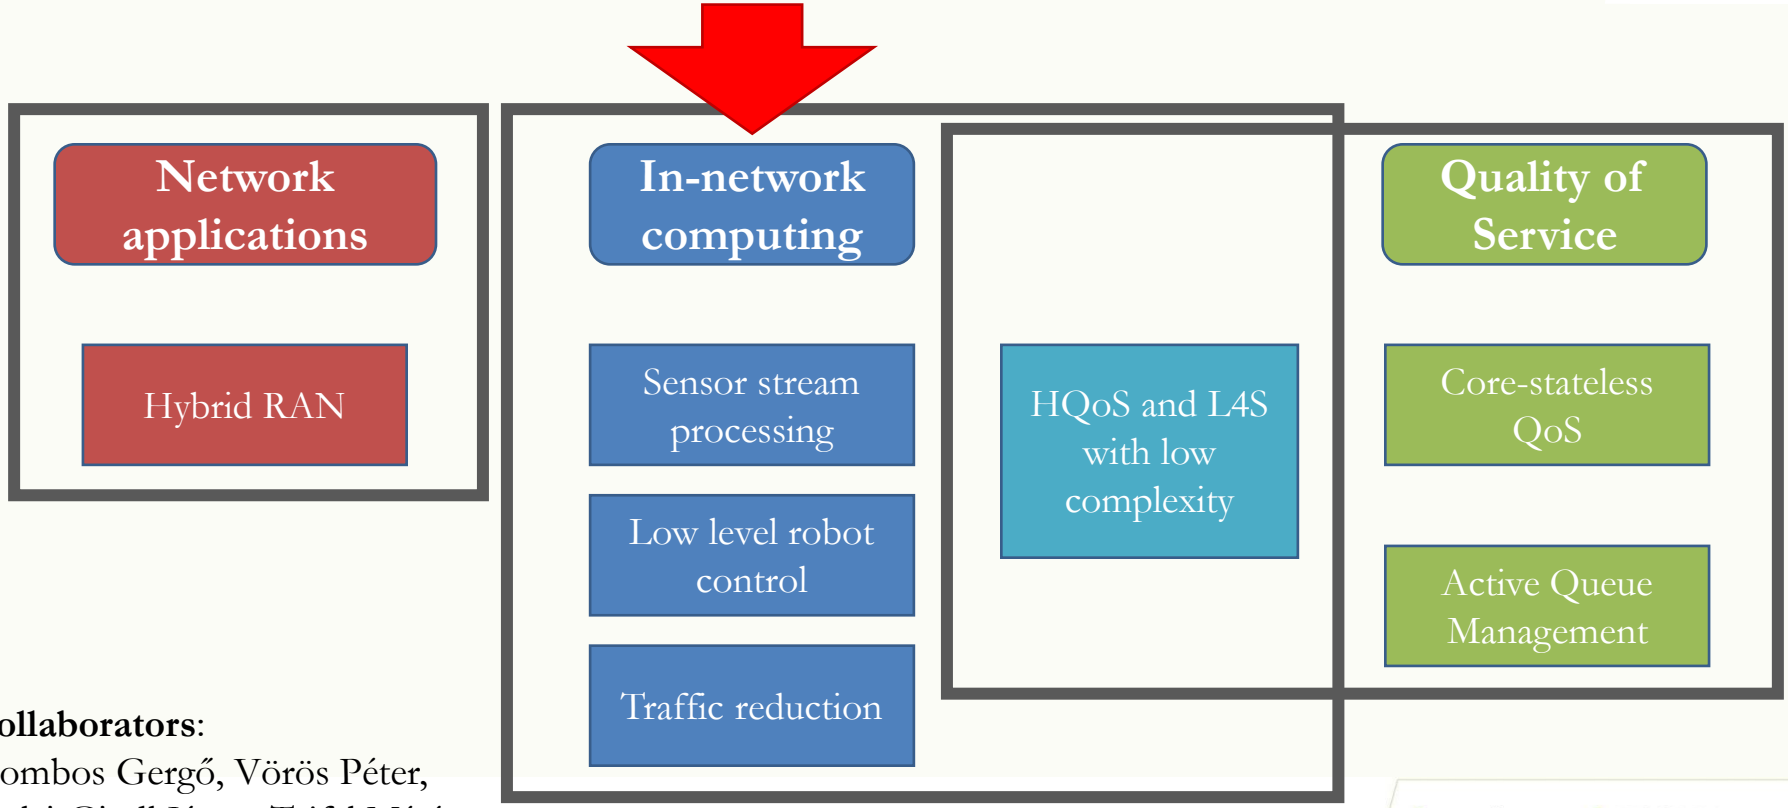

Collaborators:
Gombos Gergő, Vörös Péter, Szalai-Gindl János, Tejfel Máté, Tóthmérész Lilla, Király Tamás

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

In-network computing

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

Why in-network computing?

- **Emerging field** of networking
  - With the advent of programmable switches (BF/Intel Tofino) and P4 language

- Idea of **moving computations** from servers **to the network**
  - Enabling novel applications: caching, stream processing, query processing, load balancing, real-time control, in-network consensus, etc.

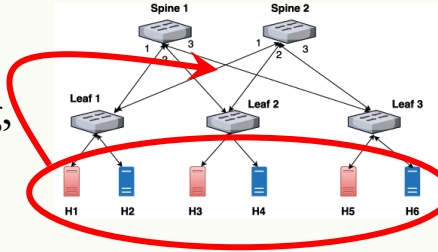

- Programmable switches are **not only** packet forwarding elements
  - Unexploited computational capacities
  - High throughput, ultra-low latency
  - Limitations
    - pipeline computing model, limited number of stages, limited memory, …

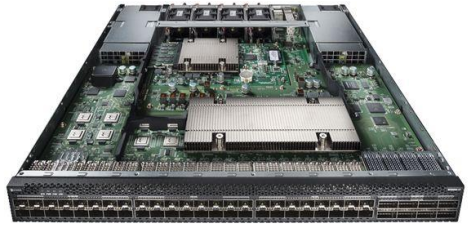

A good summary: https://www.sigarch.org/in-network-computing-draft/

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

Two recent results

In-network event detection in sensor streams

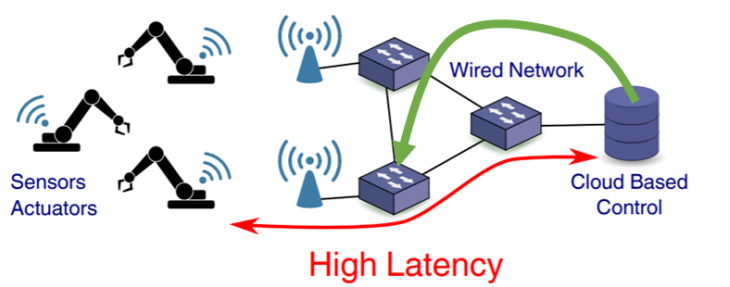

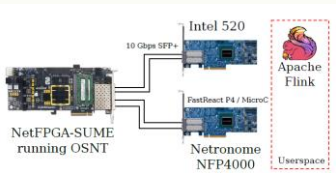

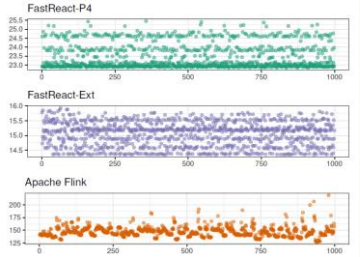

IEEE TNSM – March 2021 (Q1/D1 journal)
Joint work with Andreas Kasslerrel (KAU)

In-network robot control

Accepted @ USENIX NSDI'22 (15-16% acceptance rate)
Joint work with Ericsson Research.

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

In-network event detection with FastReact*

- **Local Decision Making** instead of centralized control
  - Early reaction reduces time required for processing
  - Reduces network data rate
  - Fewer devices that can fail

- **FastReact**
  - Implemented in P4 data plane programming language
  - Reconfigurable rules in runtime using BNF
  - Trigger local actions based on locally stored data

```
if (sensor1 > 50) && (sensor2 < 25):
    trigger_actuator(<portno>)
```

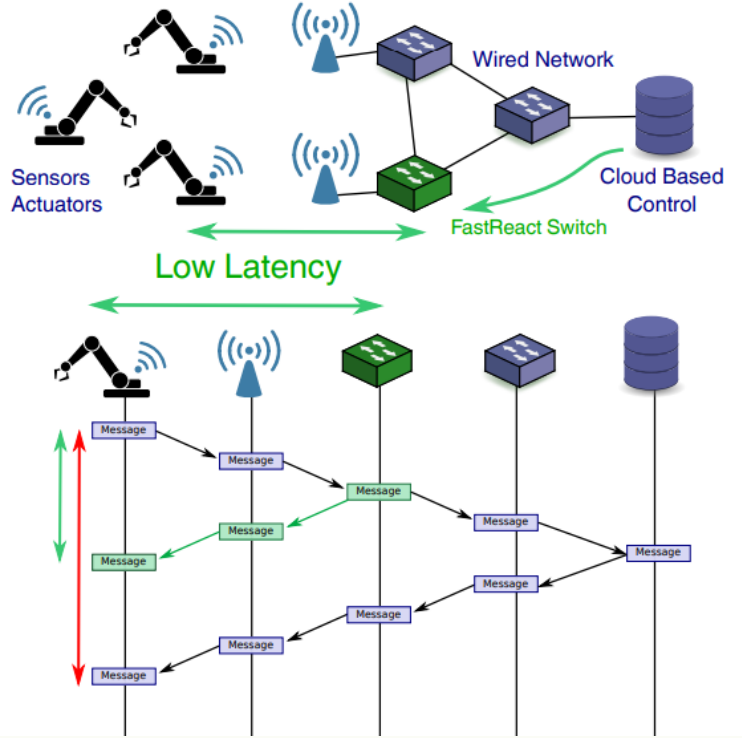

** J. Vestin, A. Kassler, S. Laki, G. Pongrácz: **Towards In-Network Event detection and Filtering for Publish/Subscribe Communication using Programmable Data Planes**, In IEEE Transactions on Network and Service Management (IEEE TNSM), Volume: 18, Issue: 1, Page(s): 415 - 428, March 2021*

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

In-Network Event Detection and Filtering for Publish/Subscribe Communication

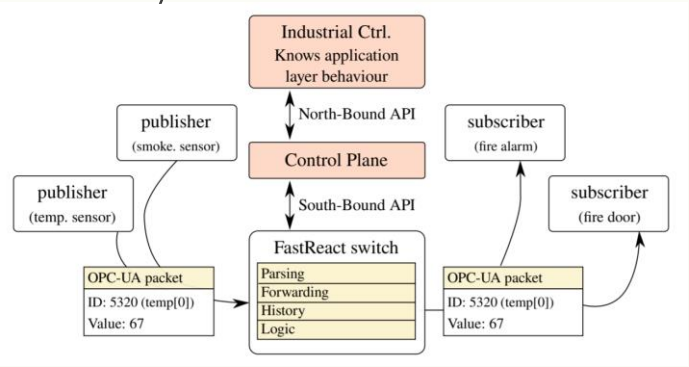

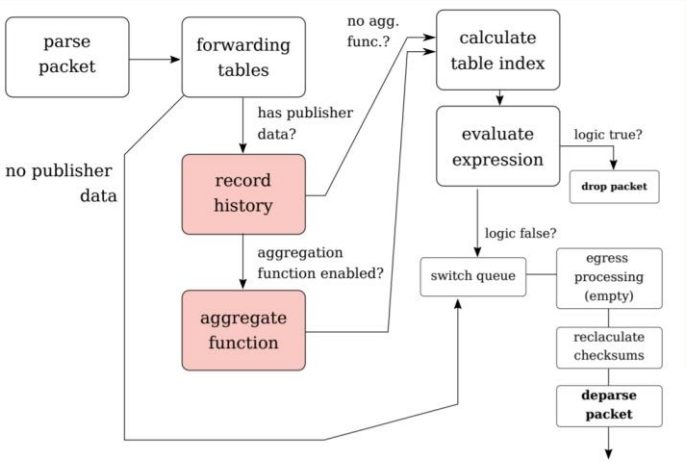

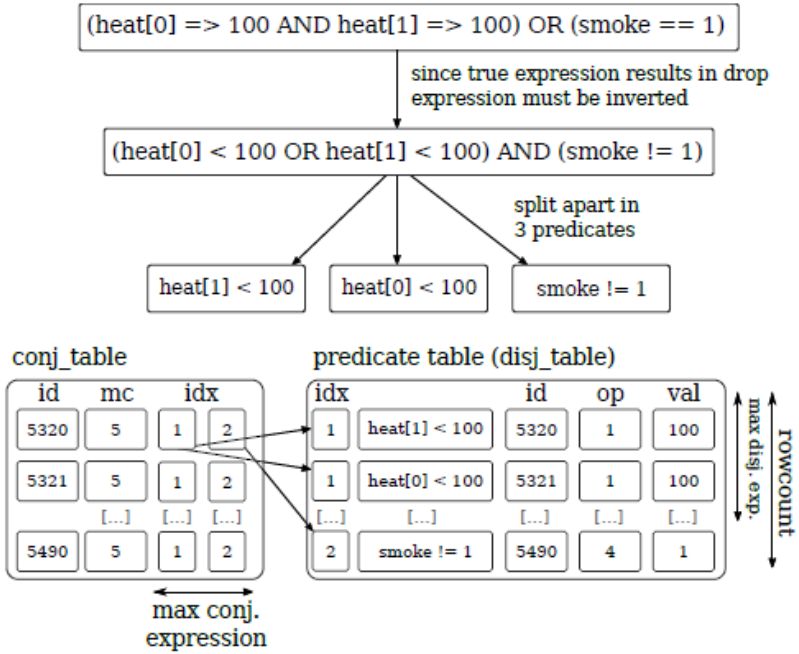

Fig. 3: FastReact table processing for the expression (heat[0] ≥ 100 ∧ heat[1] ≥ 100) ∨ (smoke = 1).

National Research, Development and Innovation Office Hungary

FINANCED FROM THE NRDI FUND

Evaluation - Latency

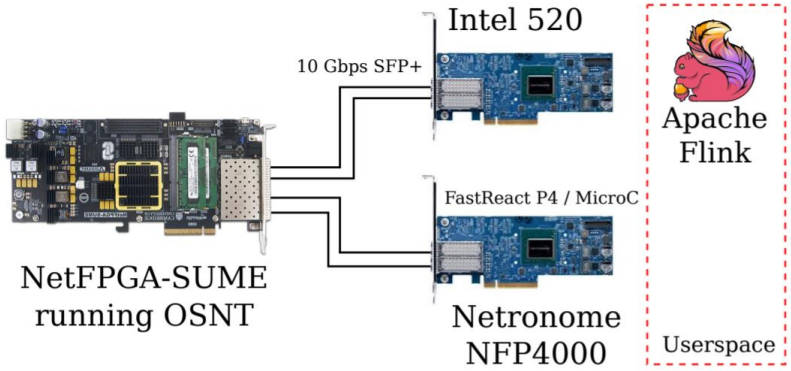

Fig. 5: Testbed setup used for evaluation

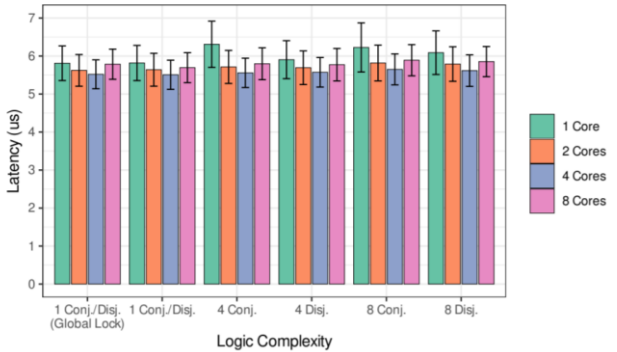

Fig. 14: Latency measurements of FastReact-PS-P4 running on the T4P4S switch, varying the disjunctive and conjunctive logic complexity. Low throughput case, preventing queue buildup.

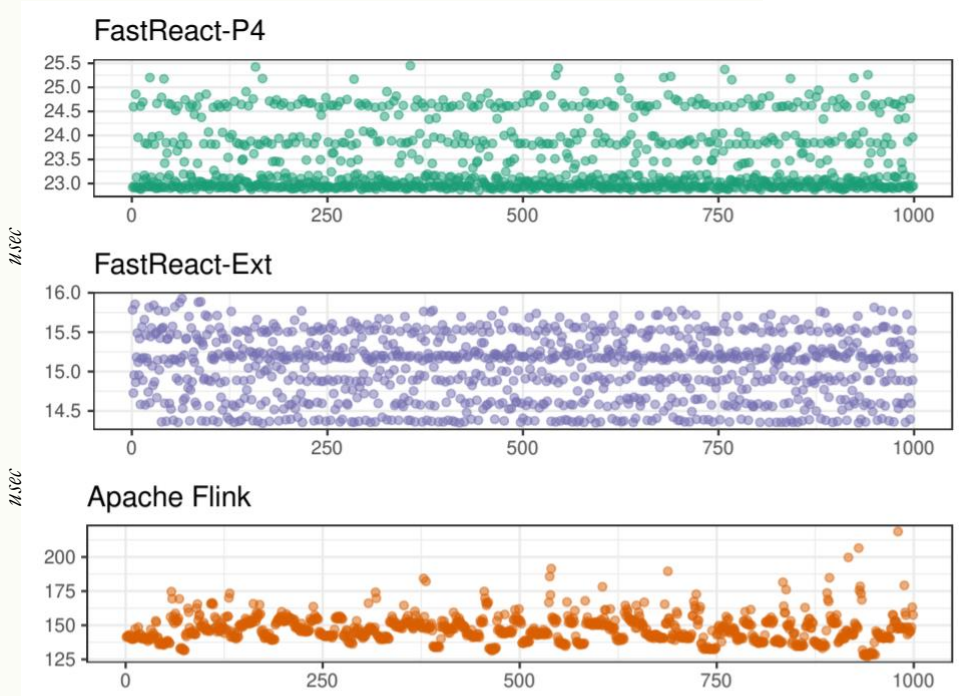

Device	Min	Mean	Max	sd	99%
FastReact-PS-P4	22.84	23.47	25.97	0.68	25.21
FastReact-PS-Ext	14.33	15.01	16.53	0.41	15.78
Apache Flink	123.88	146.69	6270.64	67.94	171.83

Latency comparison. Values are in usec.

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

Two recent results

In-network event detection in sensor streams

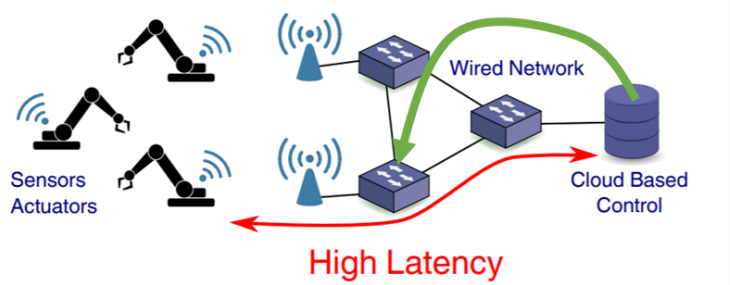

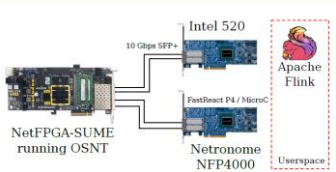

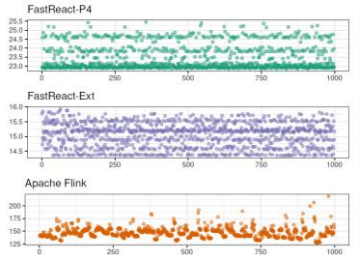

IEEE TNSM – March 2021 (Q1/D1 journal)
Joint work with Andreas Kasslerrel (KAU)

In-network robot control

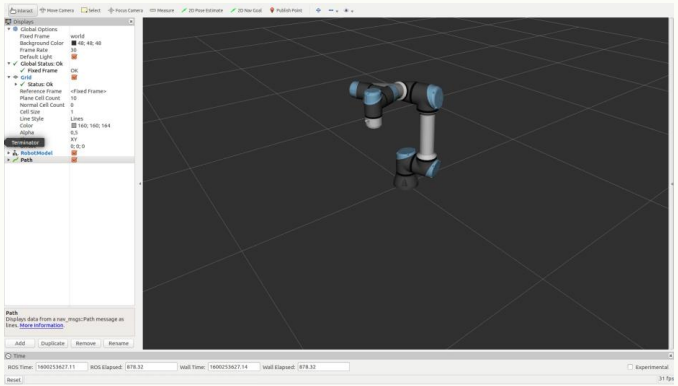

Accepted @ USENIX NSDI'22 (15-16% acceptance rate)
Joint work with Ericsson Research.

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

In-network robot control

- **Silly robots**: set of actuators with UDP-based communication
  - generate a state stream (500 Hz)
  - require velocity commands at 125 Hz
- **P4-switch**: low-level velocity control
  - **low-latency + real-time** requirements
  - responsible for calculating the velocity vectors and sending commands to the robots
- **Industrial controller**: high-level control
  - **running in VM, no real-time requirements**
  - calculating trajectories, etc.

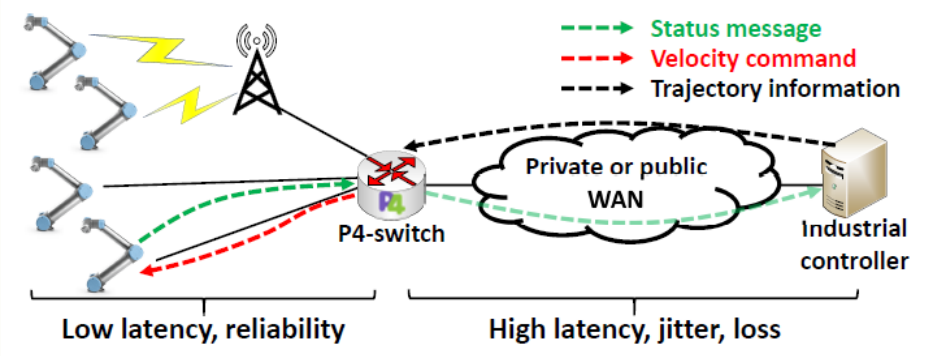

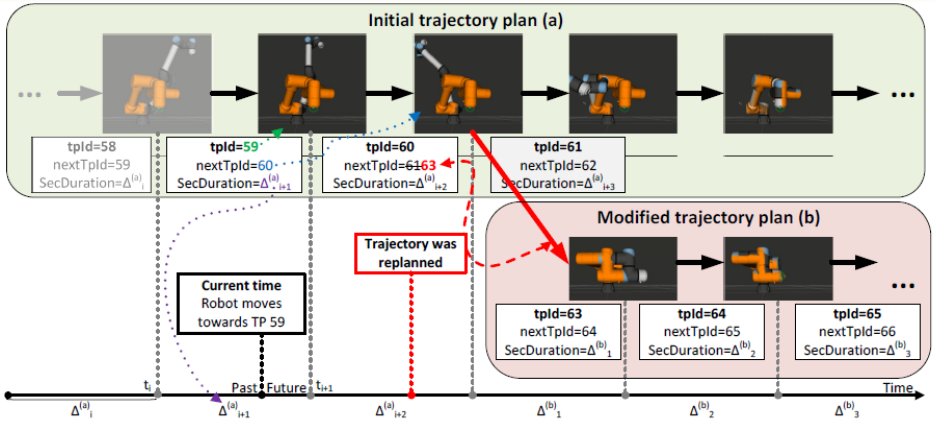

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY
PROGRAM FINANCED FROM THE NRDI FUND

In-network robot control

- Solves **synchronization** of multiple robots by design
  - Internal clock of the switch
- Good **accuracy** even with our PoC controller
  - At stop position: $0.5*10^{-7}$ rad ~ (1m long arm ~ 1-2 micrometer)
- **No strict latency requirements** for the industrial controller
  - **Enabler for cloud robotics**

- Support of **plug-and-play** deployment of new robot arms
- **Scalable** (100s of 6-DoF robot arms with a single switch)
- **Cost factor is feasable**
  - **~10000 USD switch price, 500-1000 robots: ~10-20 USD/robot**
  - Energy cost, etc.
  - DSCs may be even better for this purpose: **~2000 USD**

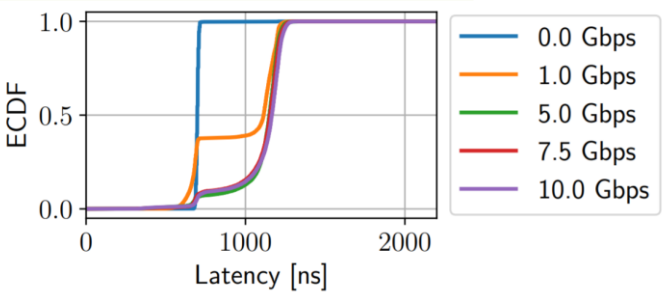

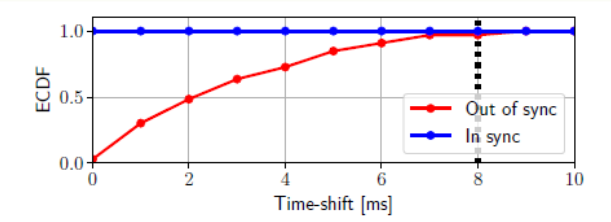

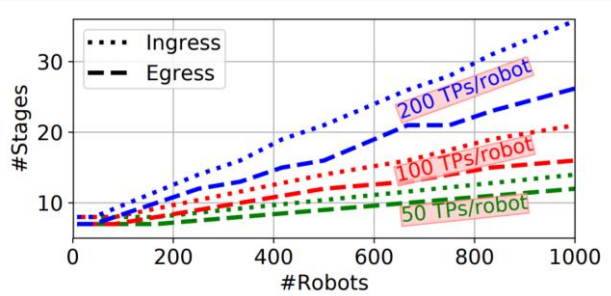

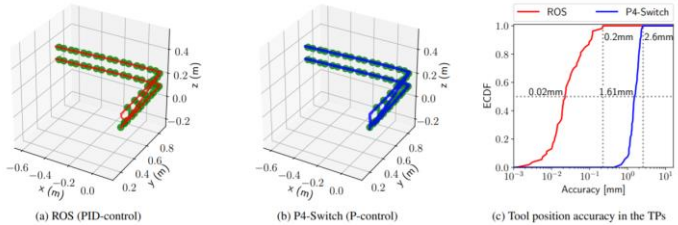

Published journal papers (Q1)
September 2020 – June 2021

- J. Vestin, A. Kassler, **S. Laki**, G. Pongrácz, Towards In-Network Event Detection and Filtering for Publish/Subscribe Communication using Programmable Data Planes, **IEEE Transactions on Network and Service Management (IEEE TNSM)** journal (Q1, IF: 5.93), Volume: 18, Issue: 1, Page(s): 415 - 428, March 2021, doi: 10.1109/TNSM.2020.3040011
- **S. Laki**, Sz. Nádas, G. Gombos, F. Fejes, P. Hudoba, Z. Turányi, Z. Kiss, Cs. Keszei, 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.
- D. Varga, J. M. Szalai-Gindl, B. Formanek, P. Vaderna, L. Dobos and **S. Laki**, Template Matching for 3D Objects in Large Point Clouds Using DBMS, in **IEEE Access** journal (Q1, IF: 3.745), vol. 9, pp. 76894-76907, 2021, doi: 10.1109/ACCESS.2021.3082848.

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

Published/Accepted conference papers
September 2020 – June 2021

- **S. Laki** et al., TBA, **19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21 – A* conf.)**, Accepted full paper
- D. A. AlWahab, G. Gombos, **S. Laki**, On a Deep Q-Network-Based Approach for Active Queue Management, **European Conference on Networks and Communications 2021 (EuCNC'21)**, 8-11 June 2021, Porto, Portugal, Virtual Conference
- Cs. Györgyi, K. Kecskeméti, P. Vörös, G. Szabó, **S. Laki,** In-network Solution for Network Traffic Reduction in Industrial Data Communication, **IEEE International Conference on Network Softwarization (IEEE NetSoft'21)**, 28 June-2 July 2021, Tokyo, Japan, Hybrid On-line Conference
- F. Fejes, Sz. Nádas, G. Gombos, **S. Laki**, A Core-Stateless L4S Scheduler for P4-enabled hardware switches with emulated HQoS [**BEST DEMO AWARD], IEEE International Conference on Computer Communications (IEEE InfoCom'21 – A* conf.)**, 10-13 May 2021, Virtual Conference [Demo paper]
- Dhulfiqar A. AlWahab and **S. Laki**, ECN-marking with CoDel and its compatibility with different TCP congestion control algorithms. In **International Conference on Advanced Science and Engineering 2020 (ICOASE'20)**, 2020, pp. 1-6, doi: 10.1109/ICOASE51841.2020.9436575.
- P. Vörös, G. Pongrácz, **S. Laki,** Towards a Hybrid Next Generation NodeB, The **3rd P4 Workshop in Europe (EUROP4 2020)** - co-located with ACM CoNext, Dec 1, 2020 - VIRTUAL CONF.

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

Other activities and results

- **Project proposals**
  - EU COST Action Proposal
    European Research Network on Programmable Networks (ProgNets)
    - As main proposer
    - Not selected
  - H2020 NGI Pointer OC-2
    - Grant for Internet renovators
    - P4EDGE project proposal on accessible P4 programmable edge
    - Selected – 200K EUR, 12 month
  - OTKA FK_21
    - Theoretical foundations of an expert system for programable data planes
    - Joint submission with Máté Tejfel

- **Two TechTalks at P4 Workshop 2021**

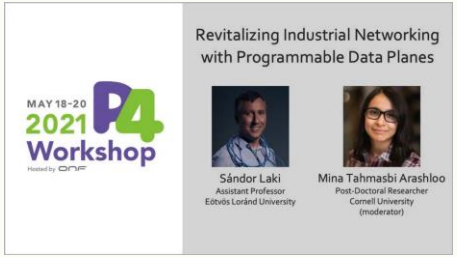

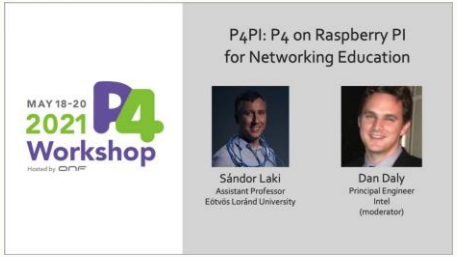

- **New project with ONF P4 Edu. WG**
  - P4PI: P4 on Raspberry Pi for Networking Education
    - Collaborators: Noa Zilberman (Oxford) and Robert Soulé (Yale)
  - **Accepted P4PI Hackathon at ACM SIGCOMM conference** (A* conf.)

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

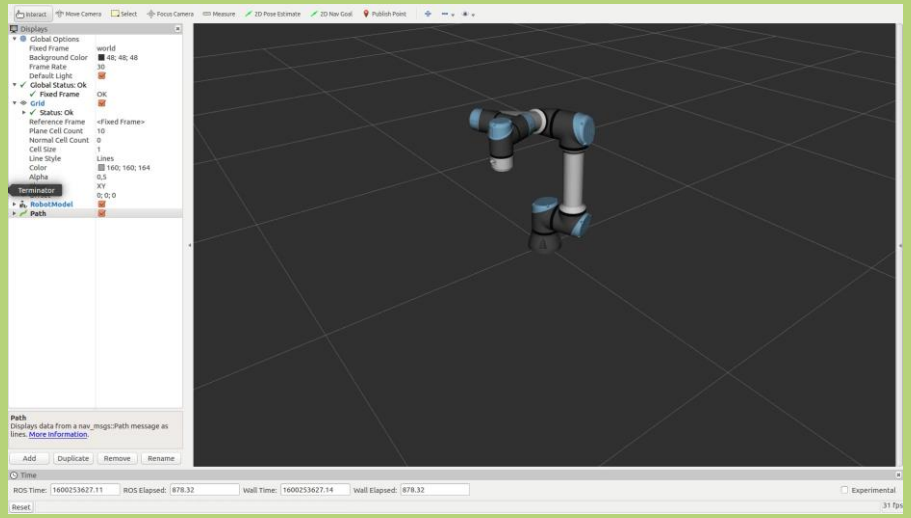

Q&A

WEB: [HTTP://LAKIS.WEB.ELTE.HU](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