NETWORK QUALITY OF SERVICE ON BIG SCALE

GERGŐ GOMBOS

PROGRAM FINANCED FROM THE NRDI FUND

Results (2022 aug. – 2023 jan.)

- **Accepted Paper**
  - **Gergő Gombos**, Dávid Kis, Lilla Tóthmérész, Tamás Király, Szilveszter Nádas, and Sándor Laki: „**Flow Fairness with Core-Stateless Resource Sharing in Arbitrary Topology**", IEEE/ACM Transactions on Networking, 2022 **(Q1)**

- **Presented Demo**
  - Dávid Kis, **Gergő Gombos**, Szilveszter Nádas, and Sándor Laki: „**Resource Sharing Beyond FQ: 35K Users at 100Gbps**", IEEE/ACM SIGCOMM Demo, 2022 **(A*)**

- **Ongoing Project**
  - Network Hierachical Quality of Service

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

Reminder: Core-Stateless Resource Sharing Control with PPV

- PPV is a Core-Stateless Resource Sharing solution
- Marking at the network edge
  - every flow has an own policy
  - the packet marked based on this policy
- AQM
  - decision based on the packet values only
  - no information about the
    - policies
    - flows
  - fast and easy implementation

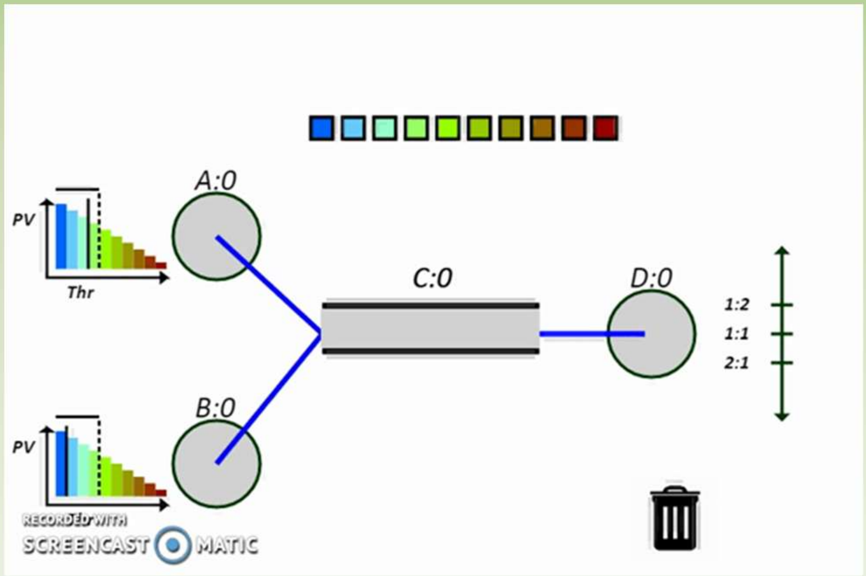

Tutorial video @ ppv.elte.hu

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

Results (2022 aug. – 2023 jan.)

- **Accepted Paper**
  - **Gergő Gombos**, Dávid Kis, Lilla Tóthmérész, Tamás Király, Szilveszter Nádas, and Sándor Laki: „**Flow Fairness with Core-Stateless Resource Sharing in Arbitrary Topology**”, IEEE/ACM Transactions on networking, 2022 (Q1)

National Research, Development and Innovation Office Hungary

PROGRAM FINANCED FROM THE NRDI FUND

PPV Core-Stateless Resource Sharing Control with PPV on large networks

- topology-zoo.org
  - real network topologies
  - nodes, edges, coordinates
  - it does not have RTTs
- Used topologies:
  - Sprint
  - Deutsche Telekom
  - Geant

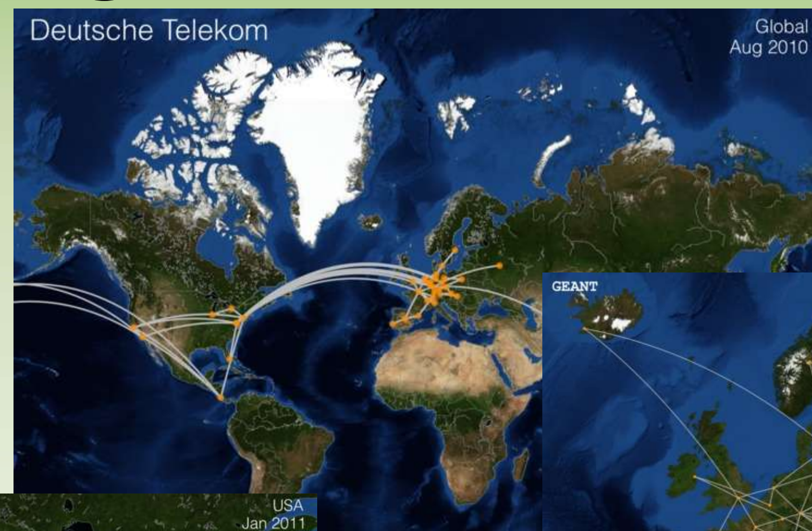

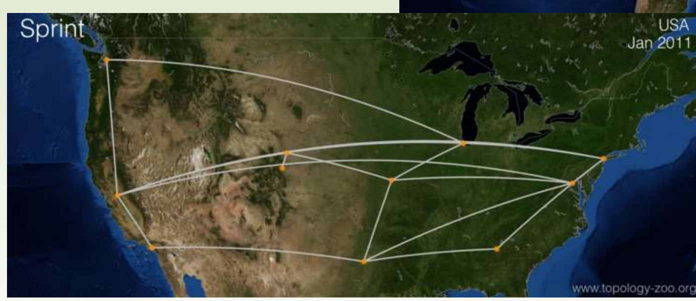

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

PPV Core-Stateless Resource Sharing Control with PPV on large networks

- Used protocols:
  - DCTCP
    - works well with small RTTs
  - NewReno
    - half the rate when dropping is
  - UDP
    - flow model simulation, no congestion control
- Simulation:
  - random $n^2/10$ flow (DT, Geant)
  - random n^2 flow (Sprint)
  - 50% Silver, 50% Gold

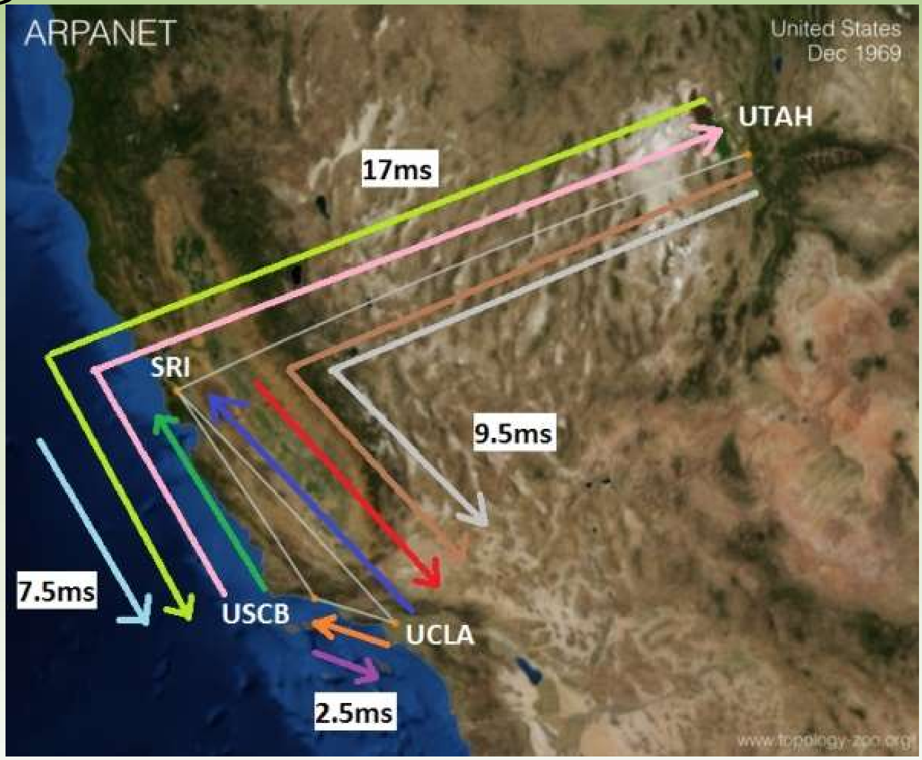

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

PPV Core-Stateless Resource Sharing Control with PPV on large networks

- 3 Silver flow on one node
- Blue lower, because it is reduce they rate earlier
- Others use the available bandwith

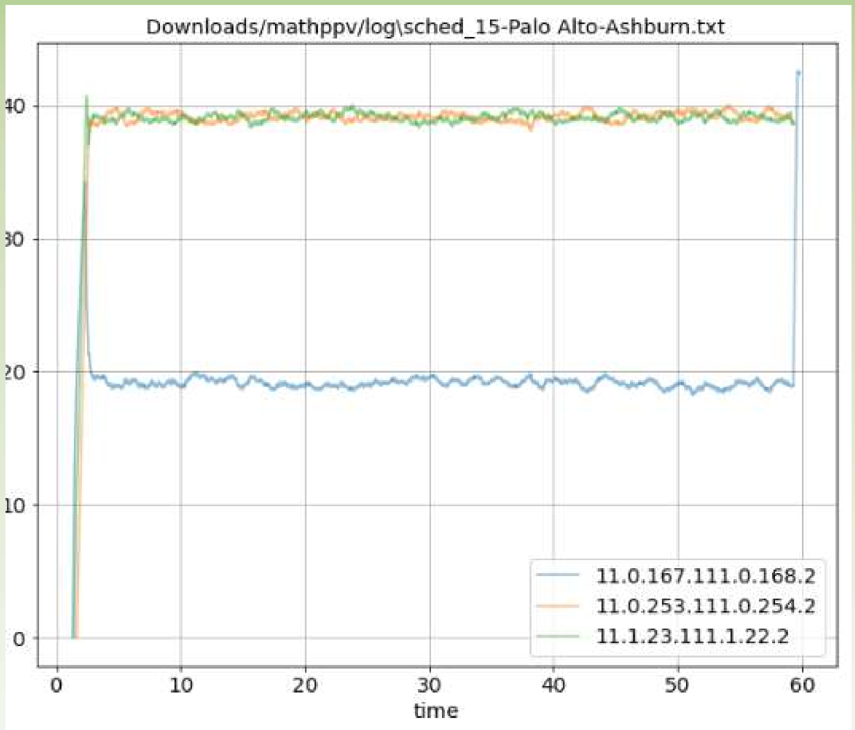

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

PPV Core-Stateless Resource Sharing Control with PPV on large networks

Deutsche Telekom

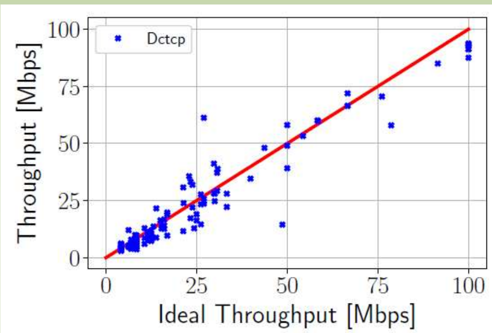

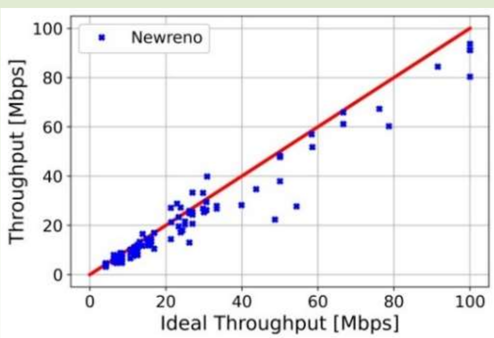

Geant

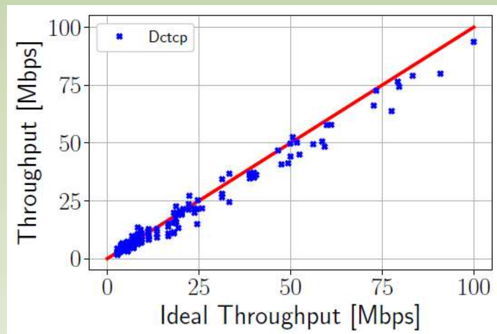

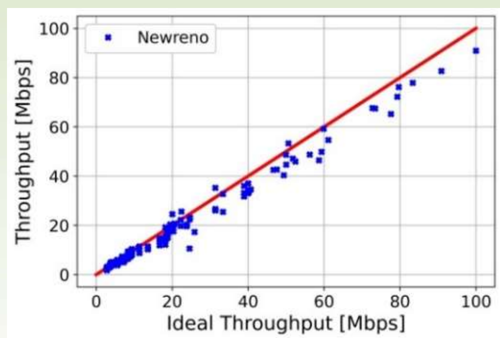

Sprint

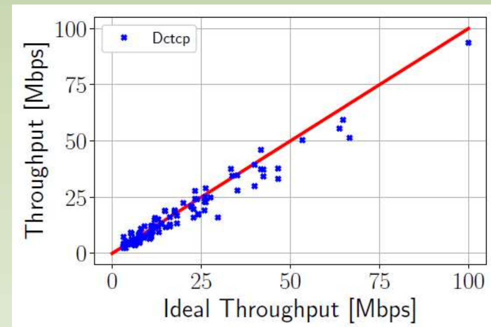

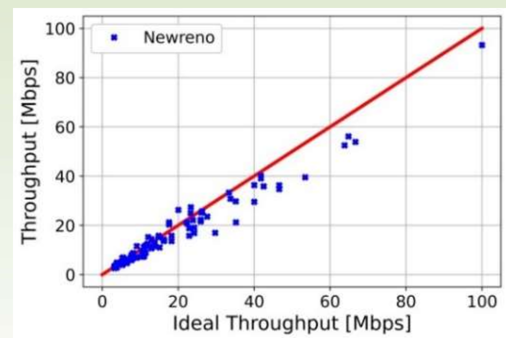

National Research, Development and Innovation Office
Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Results (2022 aug. – 2023 jan.)

- **Presented Demo**
  - Dávid Kis, **Gergő Gombos**, Szilveszter Nádas, and Sándor Laki: „**Resource Sharing Beyond FQ: 35K Users at 100Gbps**”, IEEE/ACM SIGCOMM Demo, 2022 **(A*)**

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

Testbed

communication CNL networks laboratory

ELTE EÖTVÖS LORÁND TUDOMÁNYEGYETEM

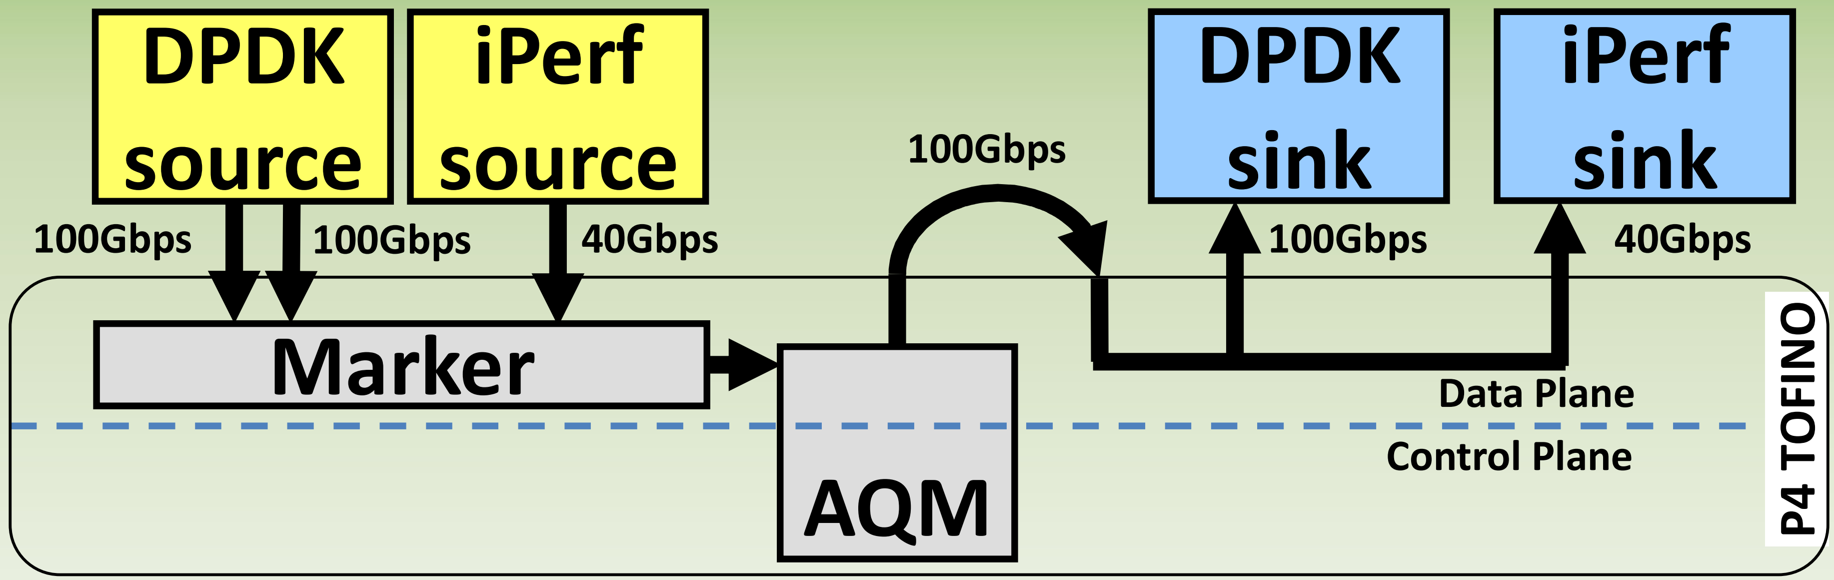

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY

PROGRAM FINANCED FROM THE NRDI FUND

Grafana dashboard

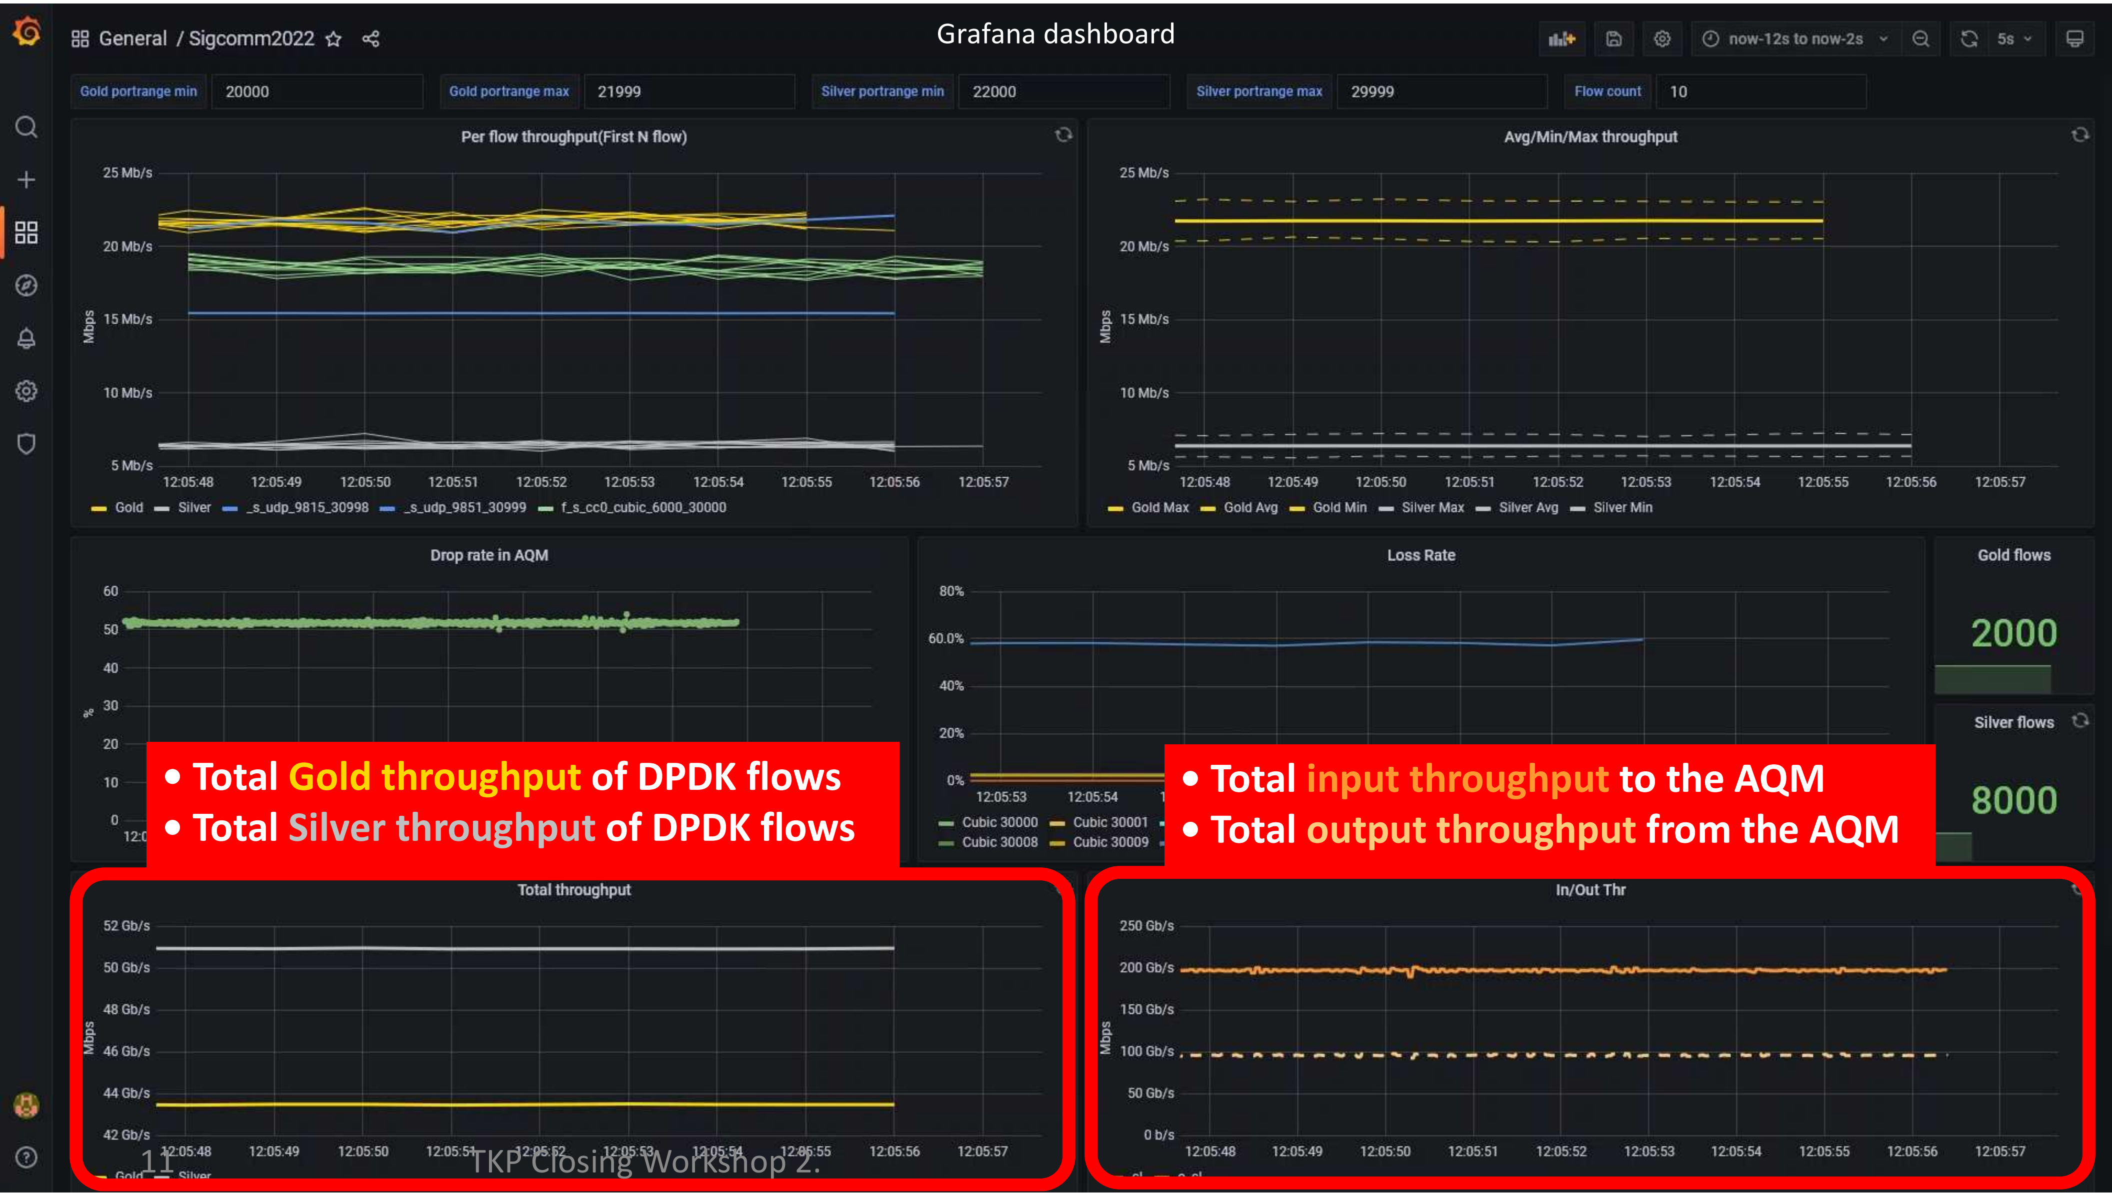

- Total **Gold throughput** of DPDK flows
- Total **Silver throughput** of DPDK flows

- Total **input throughput** to the AQM
- Total **output throughput** from the AQM

Grafana dashboard

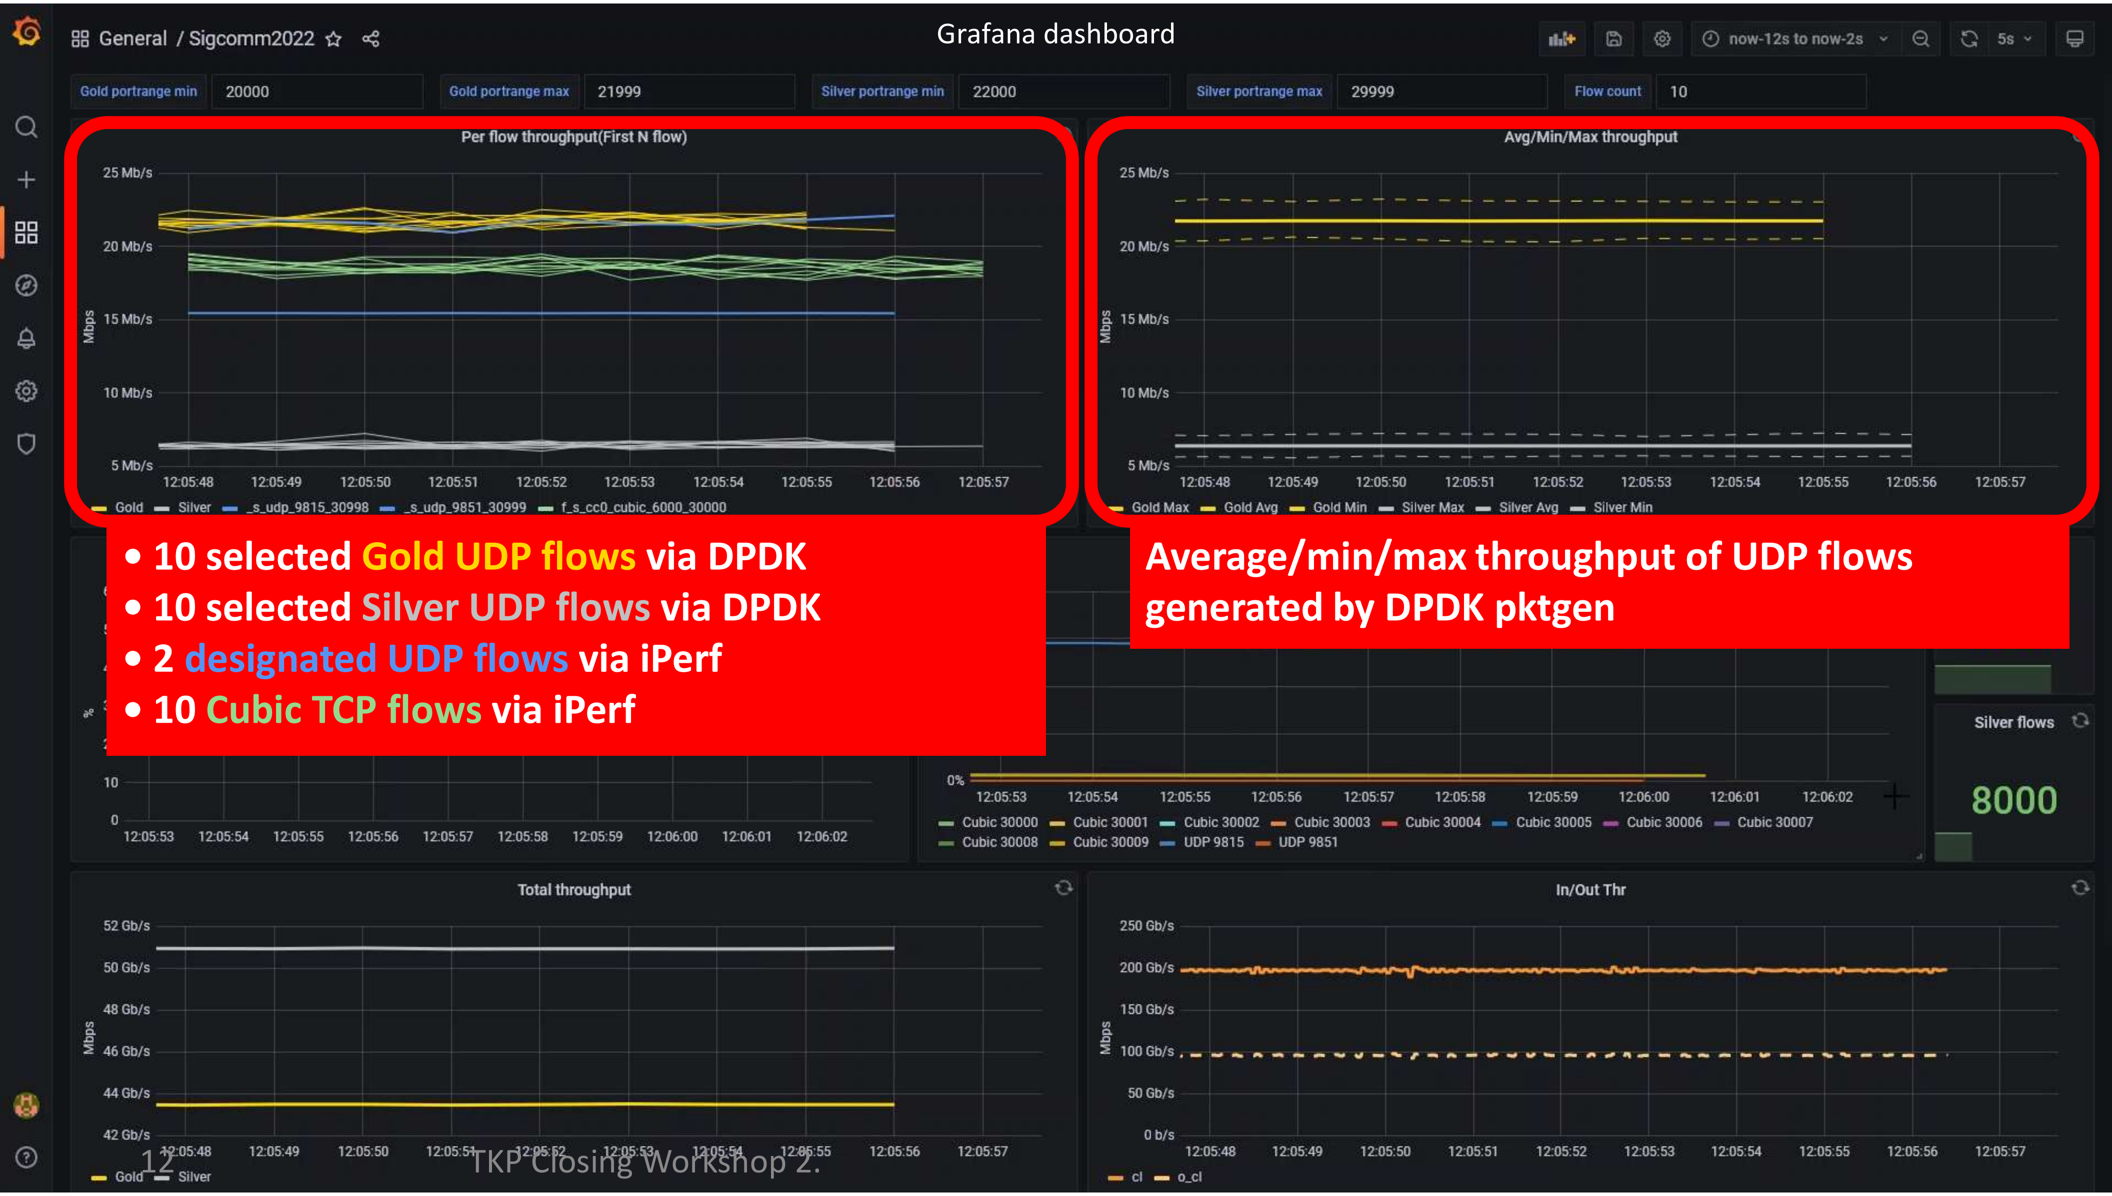

- 10 selected Gold UDP flows via DPDK
- 10 selected Silver UDP flows via DPDK
- 2 designated UDP flows via iPerf
- 10 Cubic TCP flows via iPerf

Average/min/max throughput of UDP flows generated by DPDK pktgen

35K flows
General / Sigcomm2022
2022-05-24 12:11:10 to 2022-05-24 12:11:30
Gold portrange min 20000
Gold portrange max 26999
Silver portrange min 27000
Silver portrange max 60000
Flow count 10
Per flow throughput(First N flow)
Avg/Min/Max throughput
Gold Max
Gold Avg
Gold Min
Silver Max
Silver Avg
Silver Min
Gold
Drop rate in AQM
Loss Rate
No data
Gold flows
7000
Silver flows
28000
Total throughput
Silver
In/Out Thr
cl
o_cl
Mbps

Packet Marking on P4 Tofino

communication CNL networks laboratory

ELTE EÖTVÖS LORÁND TUDOMÁNYEGYETEM

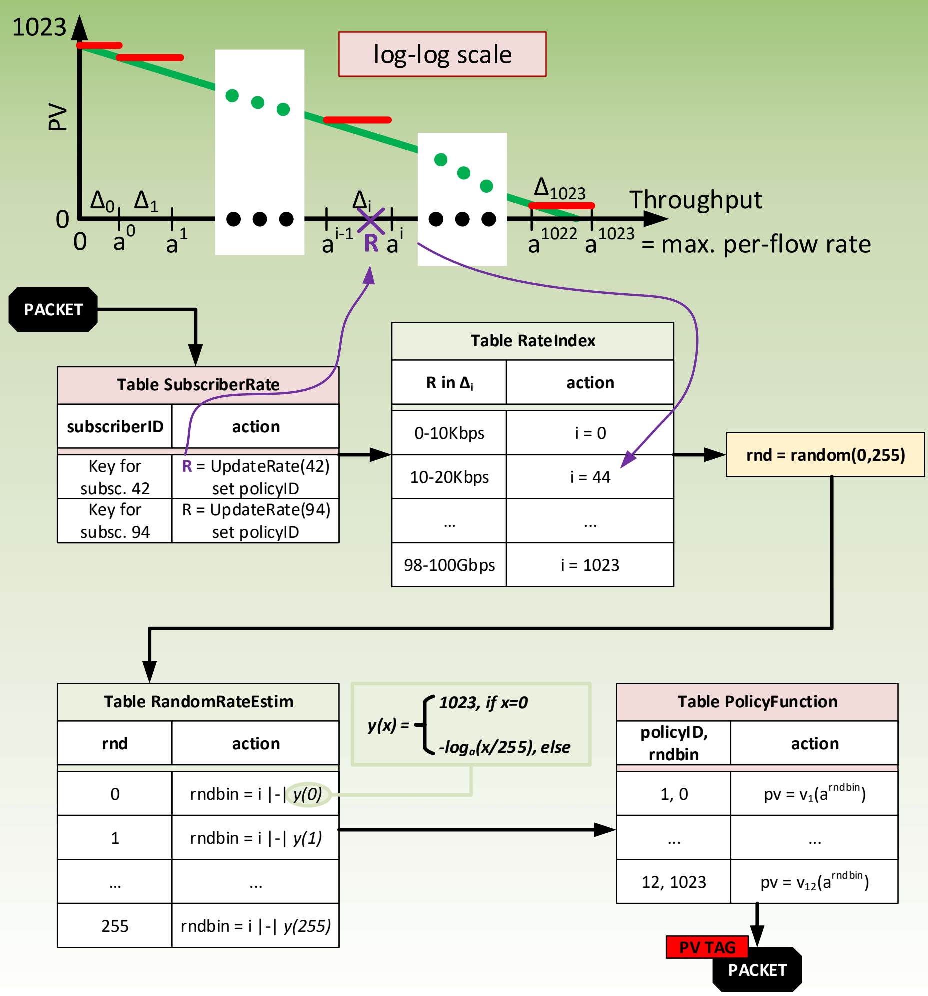

- We used the built-in rate measurement and convert it to index
- With the index and a randum number we calculate a random rate
- The policy function is quantized logarithmically and we can select the PV based on the random rate
- Data plane only implementation.
  - Policy functions are configured by the control plane
- We run 35000 marker instances in parallel

National Research, Development and Innovation Office Hungary

PROGRAM FINANCED FROM THE NRDI FUND

Results (2022 aug. – 2023 jan.)

- **Ongoing Project**
  - Network Hierachical Quality of Service

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

Mobile backhaul

- Traffic aggregates instead of flows
- Resource sharing hierarchy
- Physical network shared among multiple operators

Infrastructure-side expectation

- Utilizing resources and serving operators' needs

Operator-side expectation

- Delivering traffic between the switch site and the cells where its subscribers are located
- Getting its deserved share according to its SLA/payment

User-side expectation

- Get good service at every location

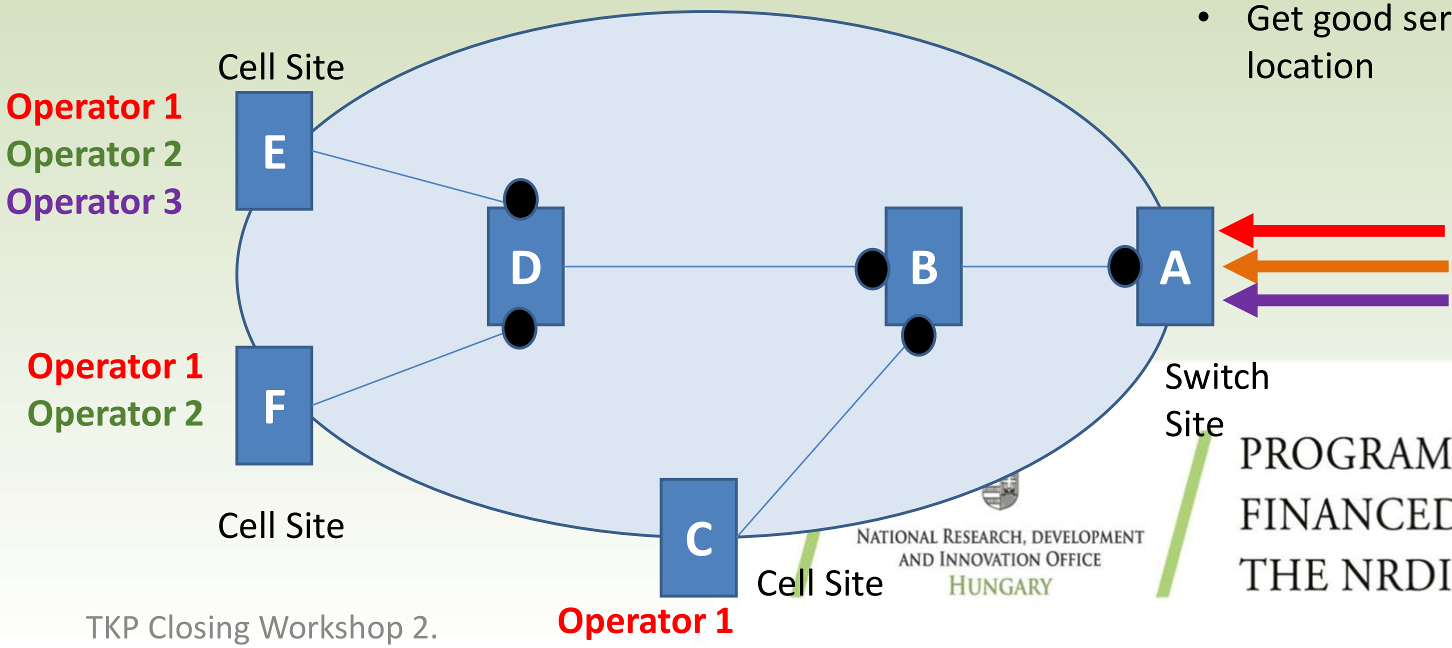

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Current approach
Link-by-link HQoS

Operator 1 = VLAN 100 -> weight 40%
Operator 2 = VLAN 200 -> weight 30%
Operator 3 = VLAN 300 -> weight 30%

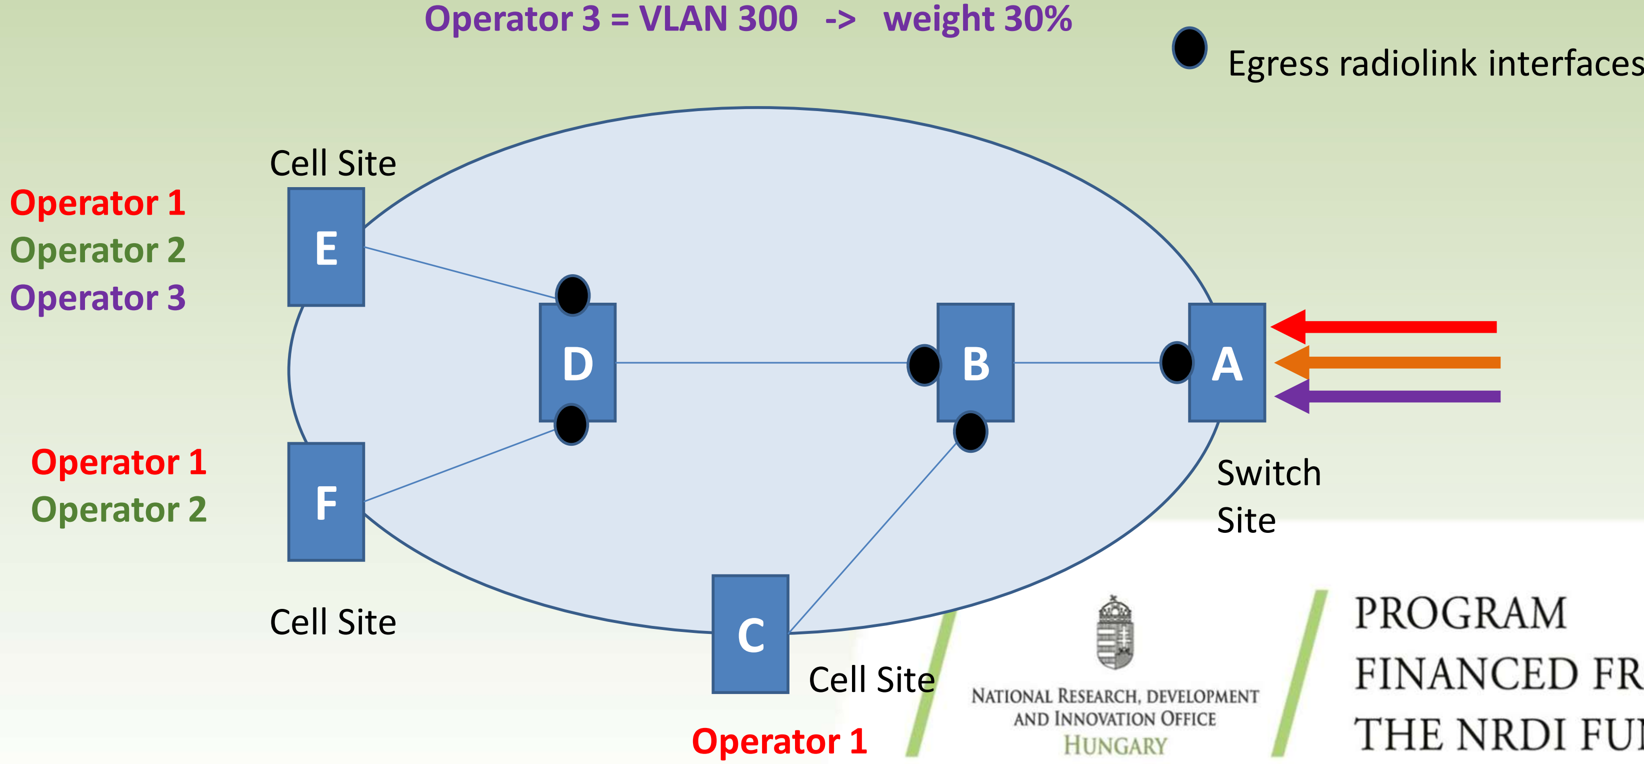

PROGRAM FINANCED FROM THE NRDI FUND

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY

Simple topology – downstream case

- Two bottleneck links
  - A-B 100Mbps
  - A-C 100Mbps

- Two operators
  - Both can send towards cell B and C
- HQoS marking at site A
  - Operator 1 and 2 has 1:1 share
  - Users within each operator are equal.

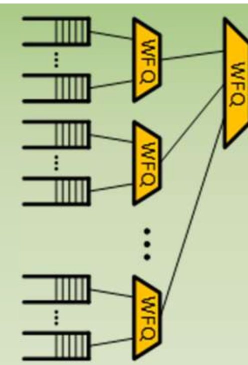

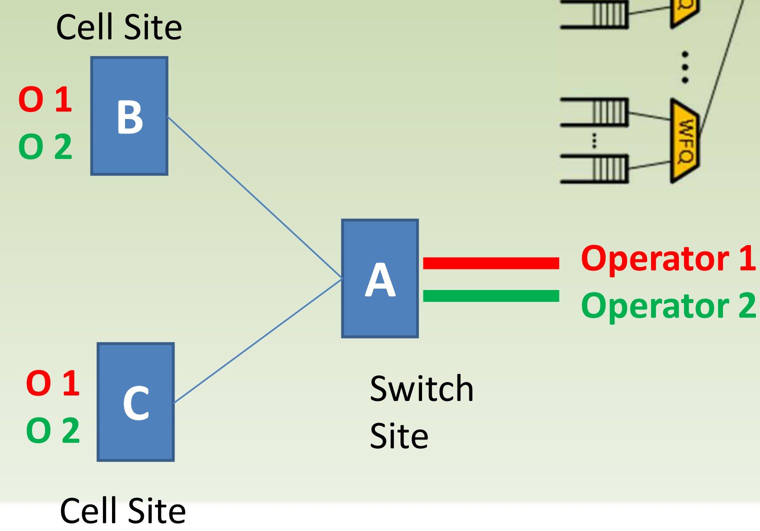

National Research, Development and Innovation Office Hungary
PROGRAM FINANCED FROM THE NRDI FUND

Link-by-link
VS
Network-wide HQoS

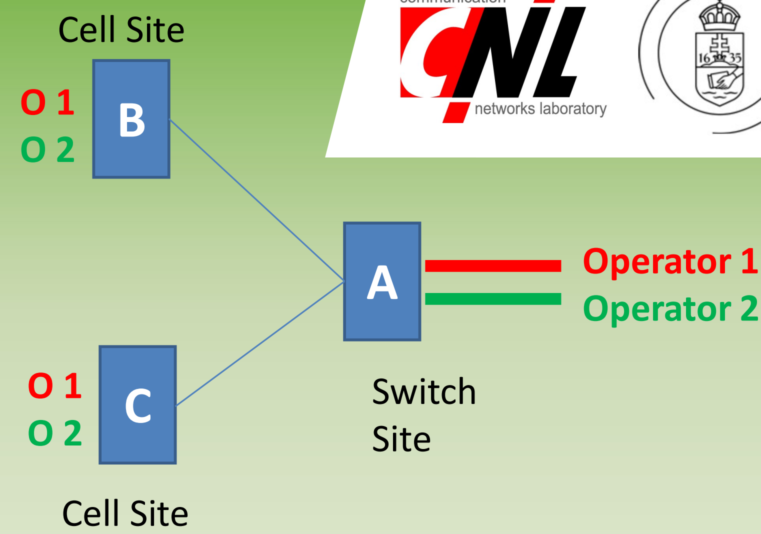

Case	Traffic				Network-Wide HQoS (NHQoS)										Link-by-link HQoS (LHQoS)									
	O1		O2		O1 throughput [Mbps]					O2 throughput [Mbps]					O1 throughput [Mbps]					O2 throughput [Mbps]				
	#users		#users		per user		total			per user		total			per user		total			per user		total		
	AB	AC	AB	AC	AB	AC	AB	AC	total	AB	AC	AB	AC	total	AB	AC	AB	AC	total	AB	AC	AB	AC	total
1	1		4		50		50		*50*	**12.5**		50		*50*	50		50		50	12.5		50		50
2	1	3	3	1	**25**	**25**	25	75	100	**25**	**25**	75	25	100	50	16.6	50	50	100	16.6	50	50	50	100
3	1	3	1	3	50	**16.6**	50	50	100	50	**16.6**	50	50	100	50	16.6	50	50	100	50	16.6	50	50	100
4	1	3	4		**20**	33.3	20	100	*120*	**20**		80		*80*	50	33.3	50	100	150	12.5		50		50
5	3	1	4		**14.3**	100	42.9	100	*142.9*	**14.3**		57.1		*57.1*	16.6	100	50	100	150	12.5		50		50
6	1	3	2	2	40	**20**	40	60	100	30	**20**	60	40	100	50	16.6	50	50	100	25	25	50	50	100
7	1	3	20	20	40	20	40	60	100	3	**2**	60	40	100	50	16.6	50	50	100	2.5	2.5	50	50	100

Table 1: Resource sharing examples. The minimum per-user throughput (marked with boldface) and the user fairness within each MNO are improved in most cases.

Case 2 „balanced asymmetry" Link-by-link HQoS

- Each link is shared according to the policy defined between the operators (1:1 in this example)
  - 100 Mbps allocated to both ops.
- The cell loads are not considered
  - The operator throughput shared evenly among them
- Each cell splits the given 50Mbps among the Ues within

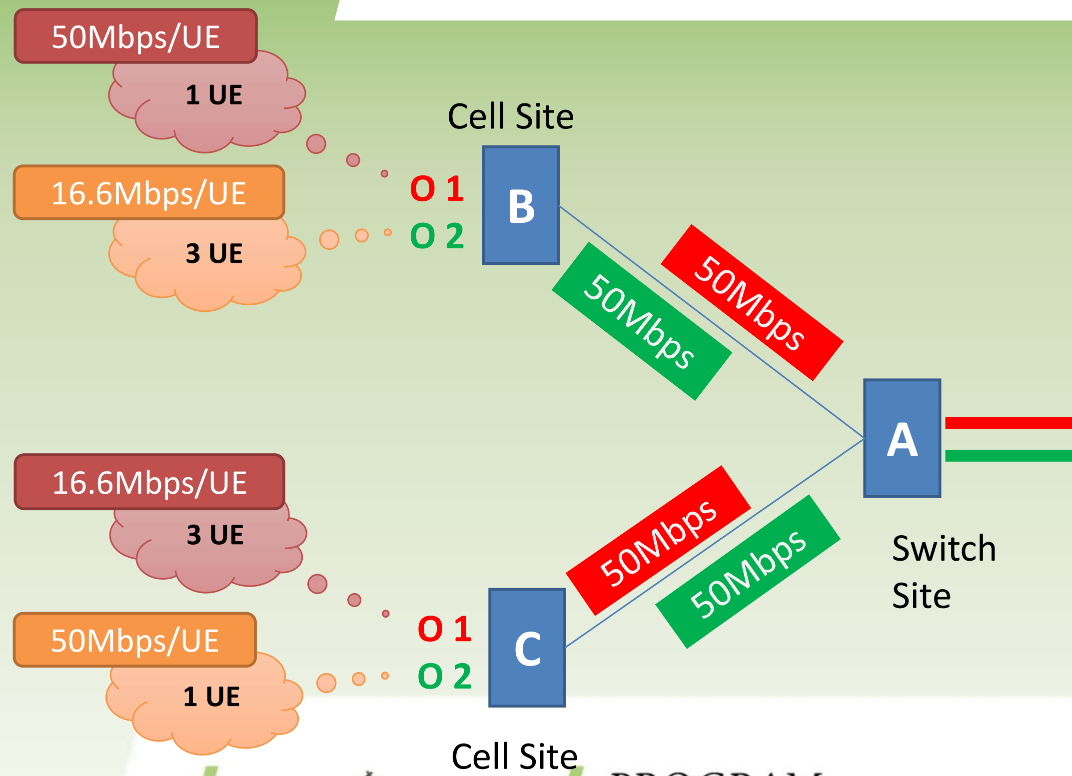

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY

PROGRAM FINANCED FROM THE NRDI FUND

Case 2 „balanced asymmetry" Network-wide HQoS

- Each link is shared proportianally
  - number of users of each operator
  - operator shares according to the policy

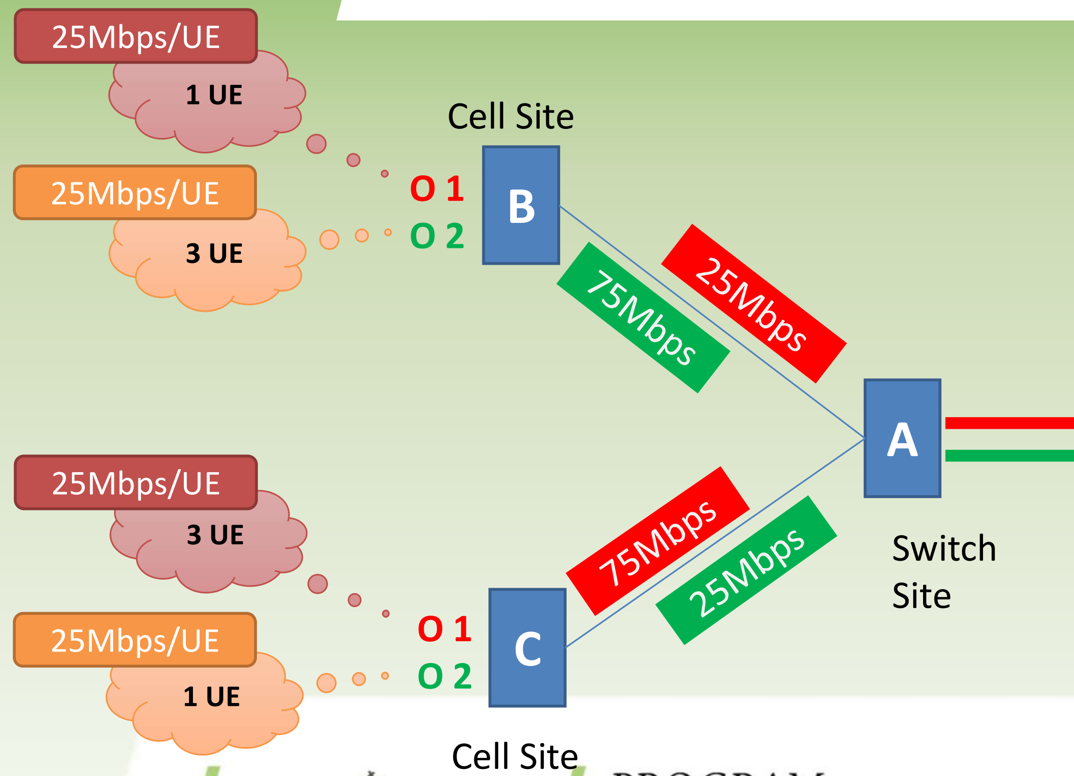

- Considers cell load

- Better fairness among UEs

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

communication CNL networks laboratory

ELTE
EÖTVÖS LORÁND
TUDOMÁNYEGYETEM

Case 4 „imbalanced asymmetry” Link-by-link HQoS

- Each link is shared according to the policy defined between the operators (1:1 in this example)
  - 100 Mbps allocated to both ops.
- Throughput allocated to UEs varies in a wide range 12.5 to 50Mbps
- Desired share between O1 and O2 is not met
  - 3:1 instead of 1:1

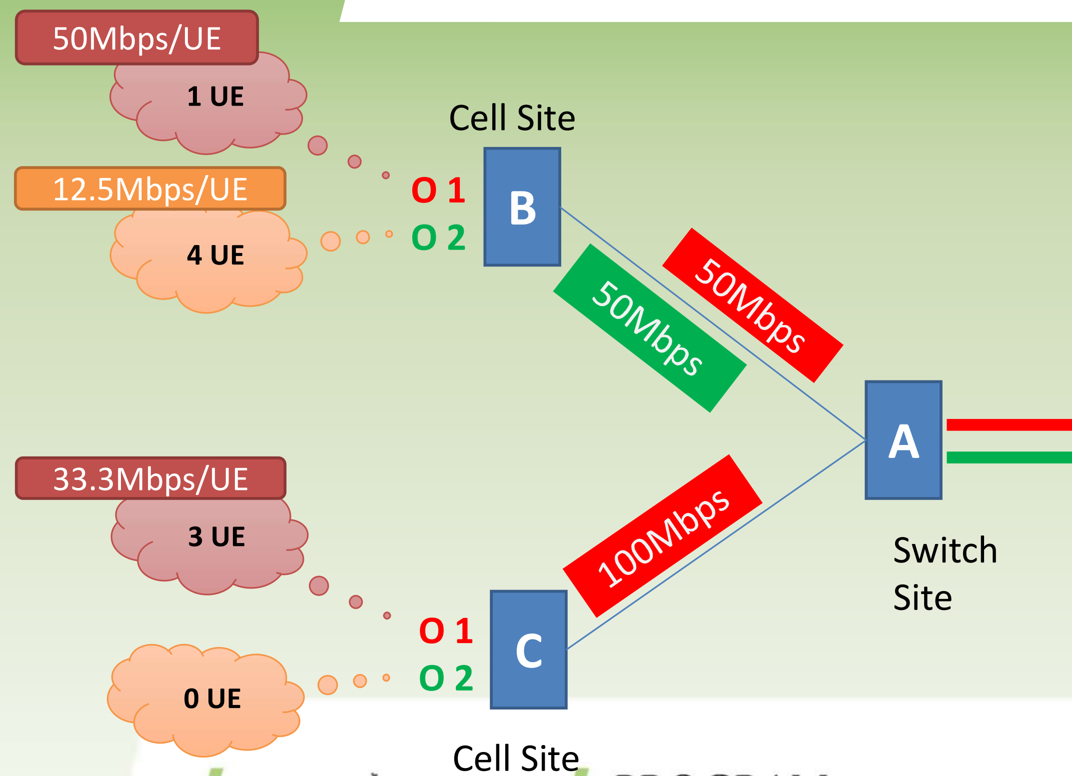

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY

PROGRAM FINANCED FROM THE NRDI FUND

Case 4 „imbalanced asymmetry" Network-wide QoS

- The most congested link is first shared proportionally
  - to the operator weigths 1:1 and the number of TAs (UEs) using the link
- On the most congested link the throughput is shared fairly among UEs in cell B
  - The other link is only used by O1 so UEs in cell C can get more without affecting UEs in cell B
- Desired share between O1 and O2 is much better than in link-by-link case
  - 12:8 instead of 1:1
  - 1:1 could only be achieved by starving a UE in cell B

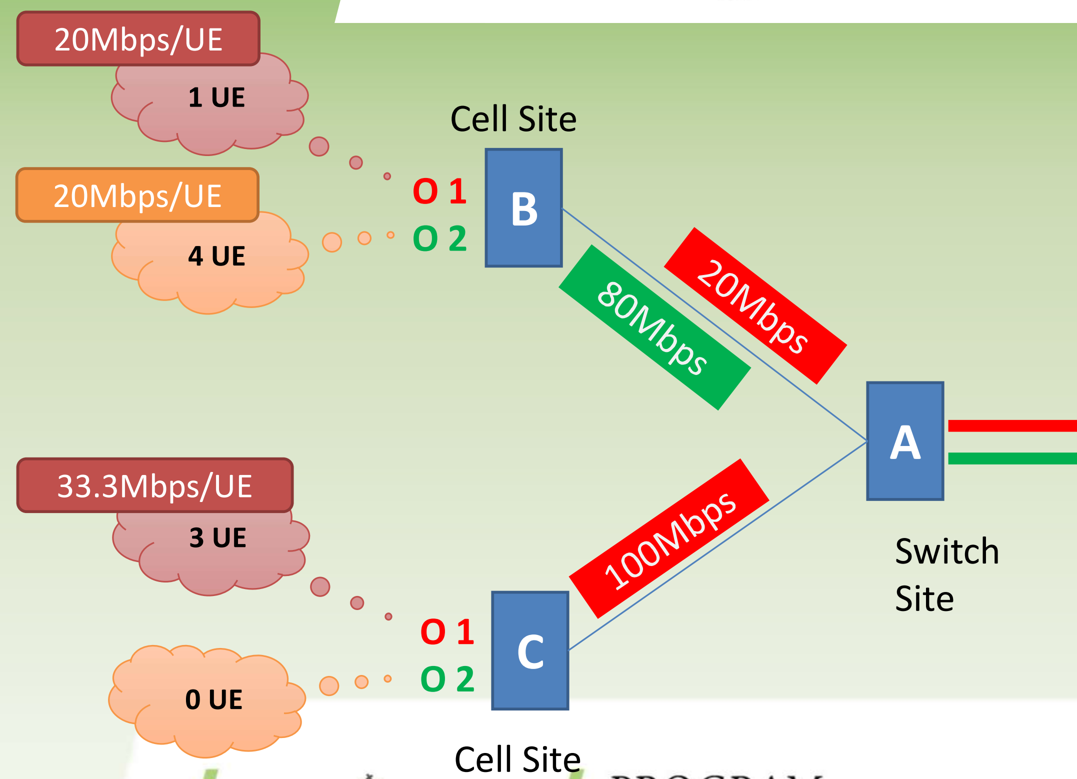

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY
PROGRAM FINANCED FROM THE NRDI FUND

Thank you!

http://ppv.elte.hu

NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION OFFICE HUNGARY

PROGRAM FINANCED FROM THE NRDI FUND