




P7 (P4 Programmable Patch Panel): an instant 100G emulated network testbed in a pizza box

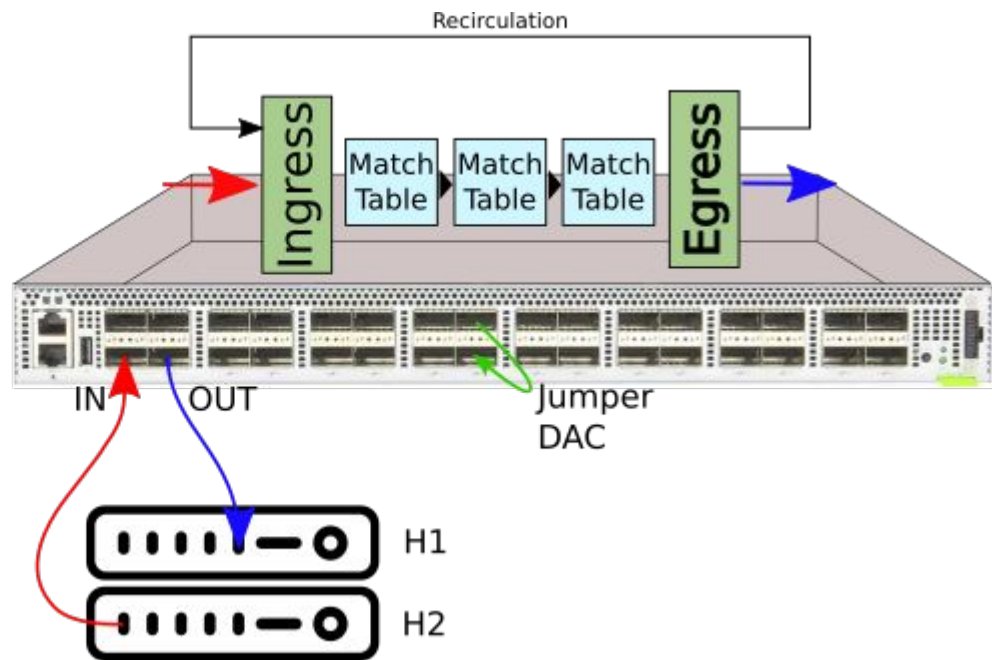
Fabricio E Rodriguez Cesen (UNICAMP),
Marcos Felipe Schwarz (RNP),
Christian Esteve Rothenberg (UNICAMP)



Agenda

- P7 overview
- Link characteristics
- Generated files
- Demo
- Topology definition and generation
- Challenges
- Future of P7


P7



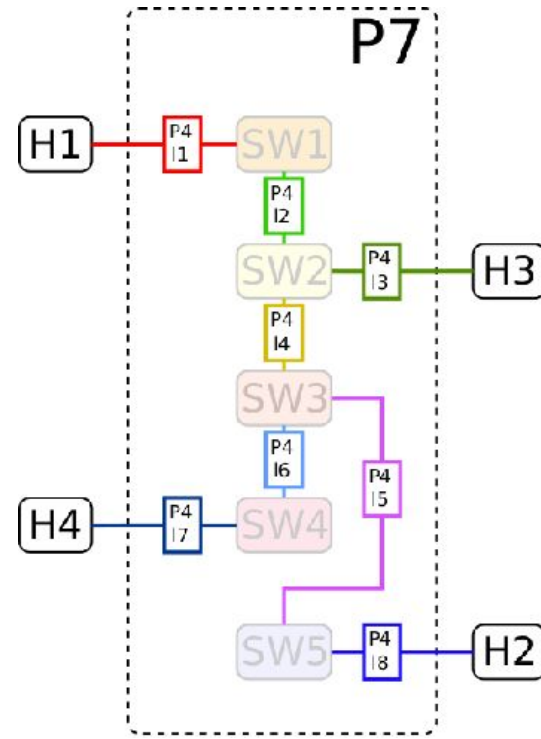
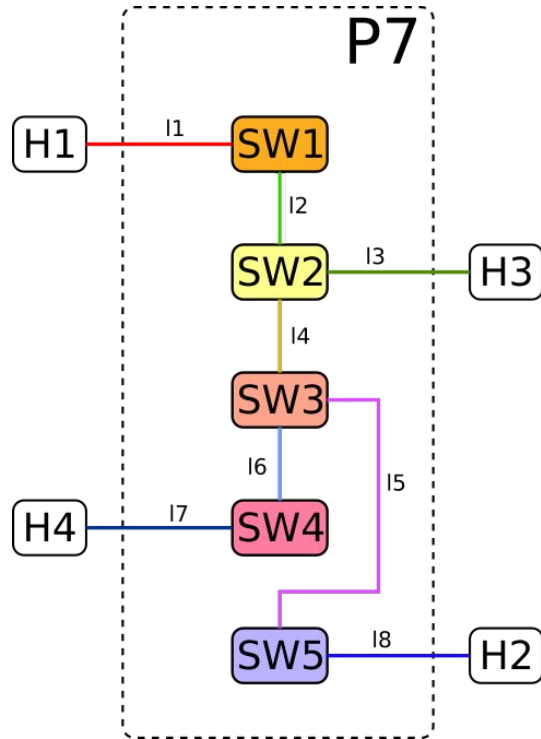
Link characteristics and implementation approaches

Link Connectivity	Jumper cabling with/without internal Tag Intern Recirculation + internal Tag
Latency [ms]	Internal timer + recirculation TM + Pipelines recirculation
Jitter [ms]	Hash to determine recirculation times Lookup table with mathematical functions
Packet loss [%]	Random function to determine the probability to discard packets Realistic packet loss model
Re-ordering	TNA TM features Targeted probabilistic packet (ID) recirculation within a flow
Bandwidth	Rate limit TNA TM feature Port configuration and mapping


Network topology





P7 topology



Generated files




Export topology from min.net



- P4 program



- Interfaces configuration
- TM



- Tables information




User needs to define the port distribution



DEMO

Link characteristics




Physical Topology




Emulated Topology

Direct link



Emulated Topology

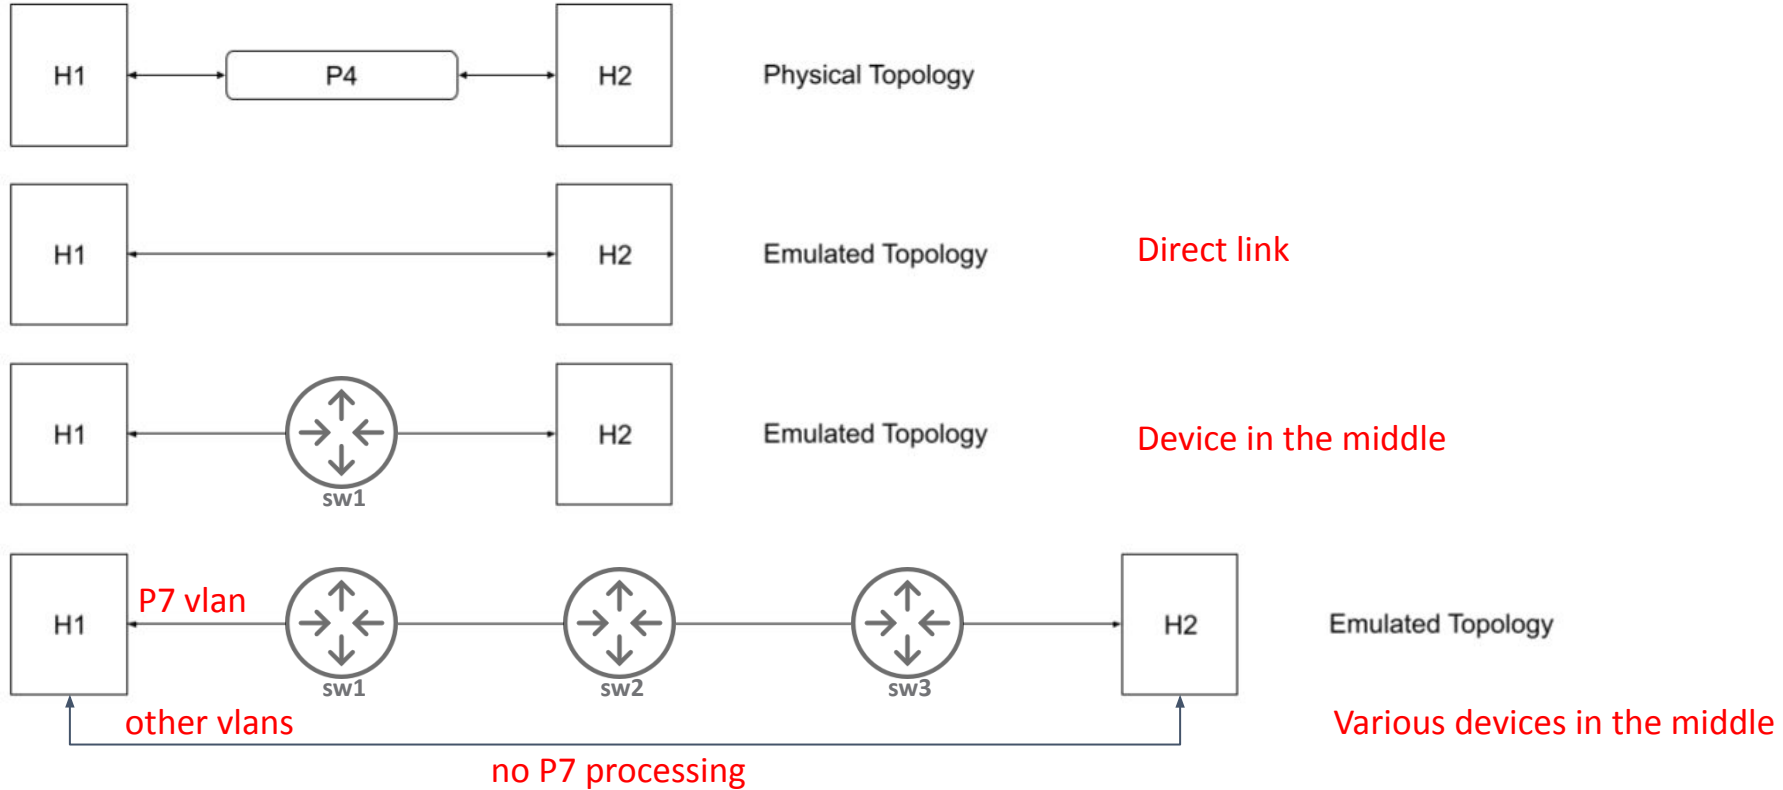
Device in the middle



Emulated Topology

Various devices in the middle

Link characteristics



Direct link

```
1 from src.data import *
2
3 topo = generator('main')
4
5 # Stratum ip:port
6 topo.addstratum("192.168.110.238:9559");
7
8 # Recirculation port default 68
9 topo.addrec_port(68)
10
11 # addhost(name,port,D_P,speed_bps,AU,FEC,vlan)
12 topo.addhost("h1",19,20,1000000000000,"False","False",1920)
13 topo.addhost("h2",20,28,1000000000000,"False","False",1920)
14
15 # addlink(node1,node2,bw,pkt_loss,latency)
16 topo.addlink("h1","h2",1000000000000,0,5)
17
18 topo.generate_chassis()
19 topo.generate_p4rt()
20 topo.generate_p4code()
21
22 topo.generate_graph()
```

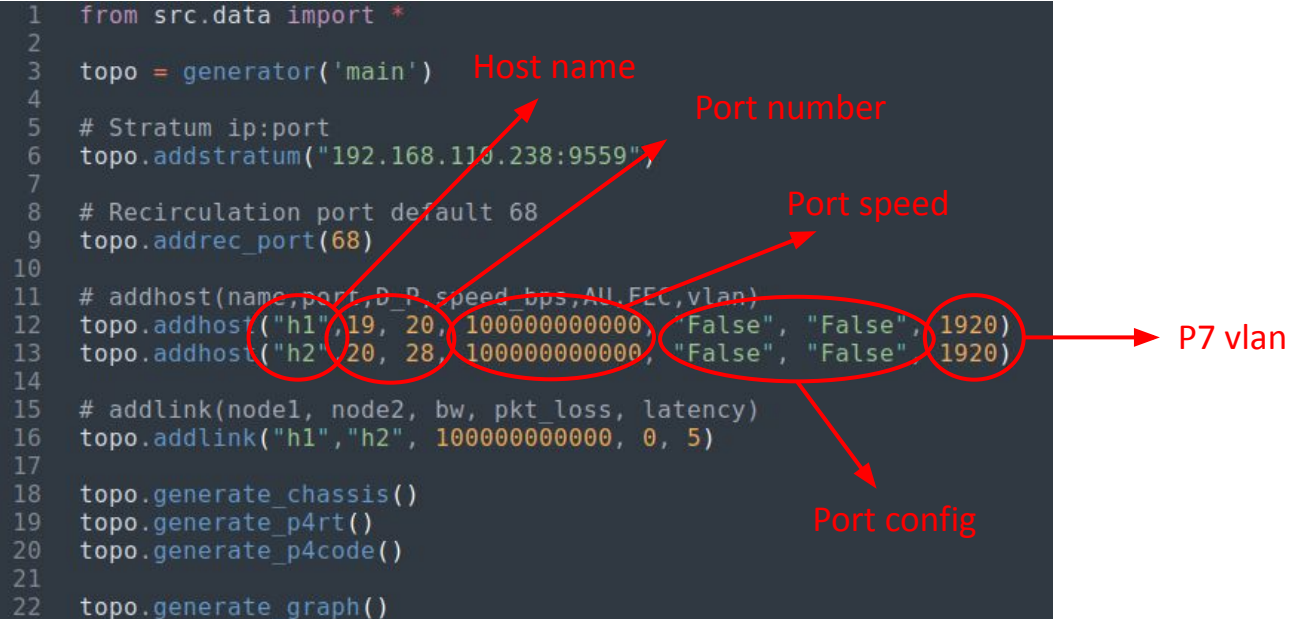
Host name

Port number

Port speed

P7 vlan

Port config

The image shows a code snippet with several red annotations. Red arrows point from text labels to specific values in the code. 'Host name' points to 'h1' and 'h2'. 'Port number' points to '19' and '20'. 'Port speed' points to '1000000000000'. 'P7 vlan' points to '1920'. 'Port config' points to the 'False' values. Red circles highlight the host names and the '1920' values in the addhost calls.

Direct link

```
1 from src.data import *
2
3 topo = generator('main') node 1
4
5 # Stratum ip:port
6 topo.addstratum("192.168.110.238:9559")
7
8 # Recirculation port default 68
9 topo.addrec_port(68)
10
11 # addhost(name,port,D_P,speed_bps,AU,FEC,vlan)
12 topo.addhost("h1",19, 20, 1000000000000, "False", "False", 1920)
13 topo.addhost("h2",20, 28, 1000000000000, "False", "False", 1920)
14
15 # addlink(node1,node2, bw, pkt_loss, latency)
16 topo.addlink("h1","h2",1000000000000,0,5)
17
18 topo.generate_chassis()
19 topo.generate_p4rt()
20 topo.generate_p4code()
21
22 topo.generate_graph()
```

node 1

node 2

link bw

latency


pkt loss

Device in the middle

```
1  from src.data import *
2
3  topo = generator('main')
4
5  # Stratum ip:port
6  topo.addstratum("192.168.110.238:9559")
7
8  # Recirculation port default 68
9  topo.addrec_port(68)
10
11 # addswitch(name)
12 topo.addswitch("sw1")
13
14 # addhost(name,port,D_P,speed_bps,AU,FEC,vlan)
15 # include the link configuration
16 topo.addhost("h1",19, 20, 100000000000, "False", "False", 1920)
17 topo.addhost("h2",20, 28, 100000000000, "False", "False", 1920)
18
19 # addlink(node1, node2, bw, pkt_loss, latency)
20 topo.addlink("h1","sw1", 100000000000, 0, 5)
21 topo.addlink("sw1","h2", 100000000000, 0, 5)
22
23 topo.generate_chassis()
24 topo.generate_p4rt()
25 topo.generate_p4code()
26
27 topo.generate_graph()
```

switch name

links



https://drive.google.com/file/d/1Hop_zgK-yIvP3MwUjHpMdhiqVvXNpiwY/view?usp=sharing

Future of P7

- Address open challenges
 - Topology Size
 - Buffers consumption
- New features
 - Packet generation for background/congestion traffic
 - In-band Network Telemetry (INT)
 - Dynamic link behaviors (e.g., pkt loss patterns)
 - Trace base link characteristics
- Open source repository and community
- Embed into disaggregated network testbed initiatives

P7 repository

Available soon at:

<https://github.com/intrig-unicamp/p7>

References / Related Work

- Mininet <http://mininet.org/>
- TurboNet: Faithfully Emulating Networks with Programmable Switches. [Link](#) [Code](#)
- CrystalNet: Faithfully Emulating Large Production Networks. [Link](#)
- BNV: Enabling Scalable Network Experimentation through Bare-metal Network Virtualization. [Link](#)



Thank You

Fabricio Rodriguez

frodri@dca.fee.unicamp.br

<https://intrig.dca.fee.unicamp.br>