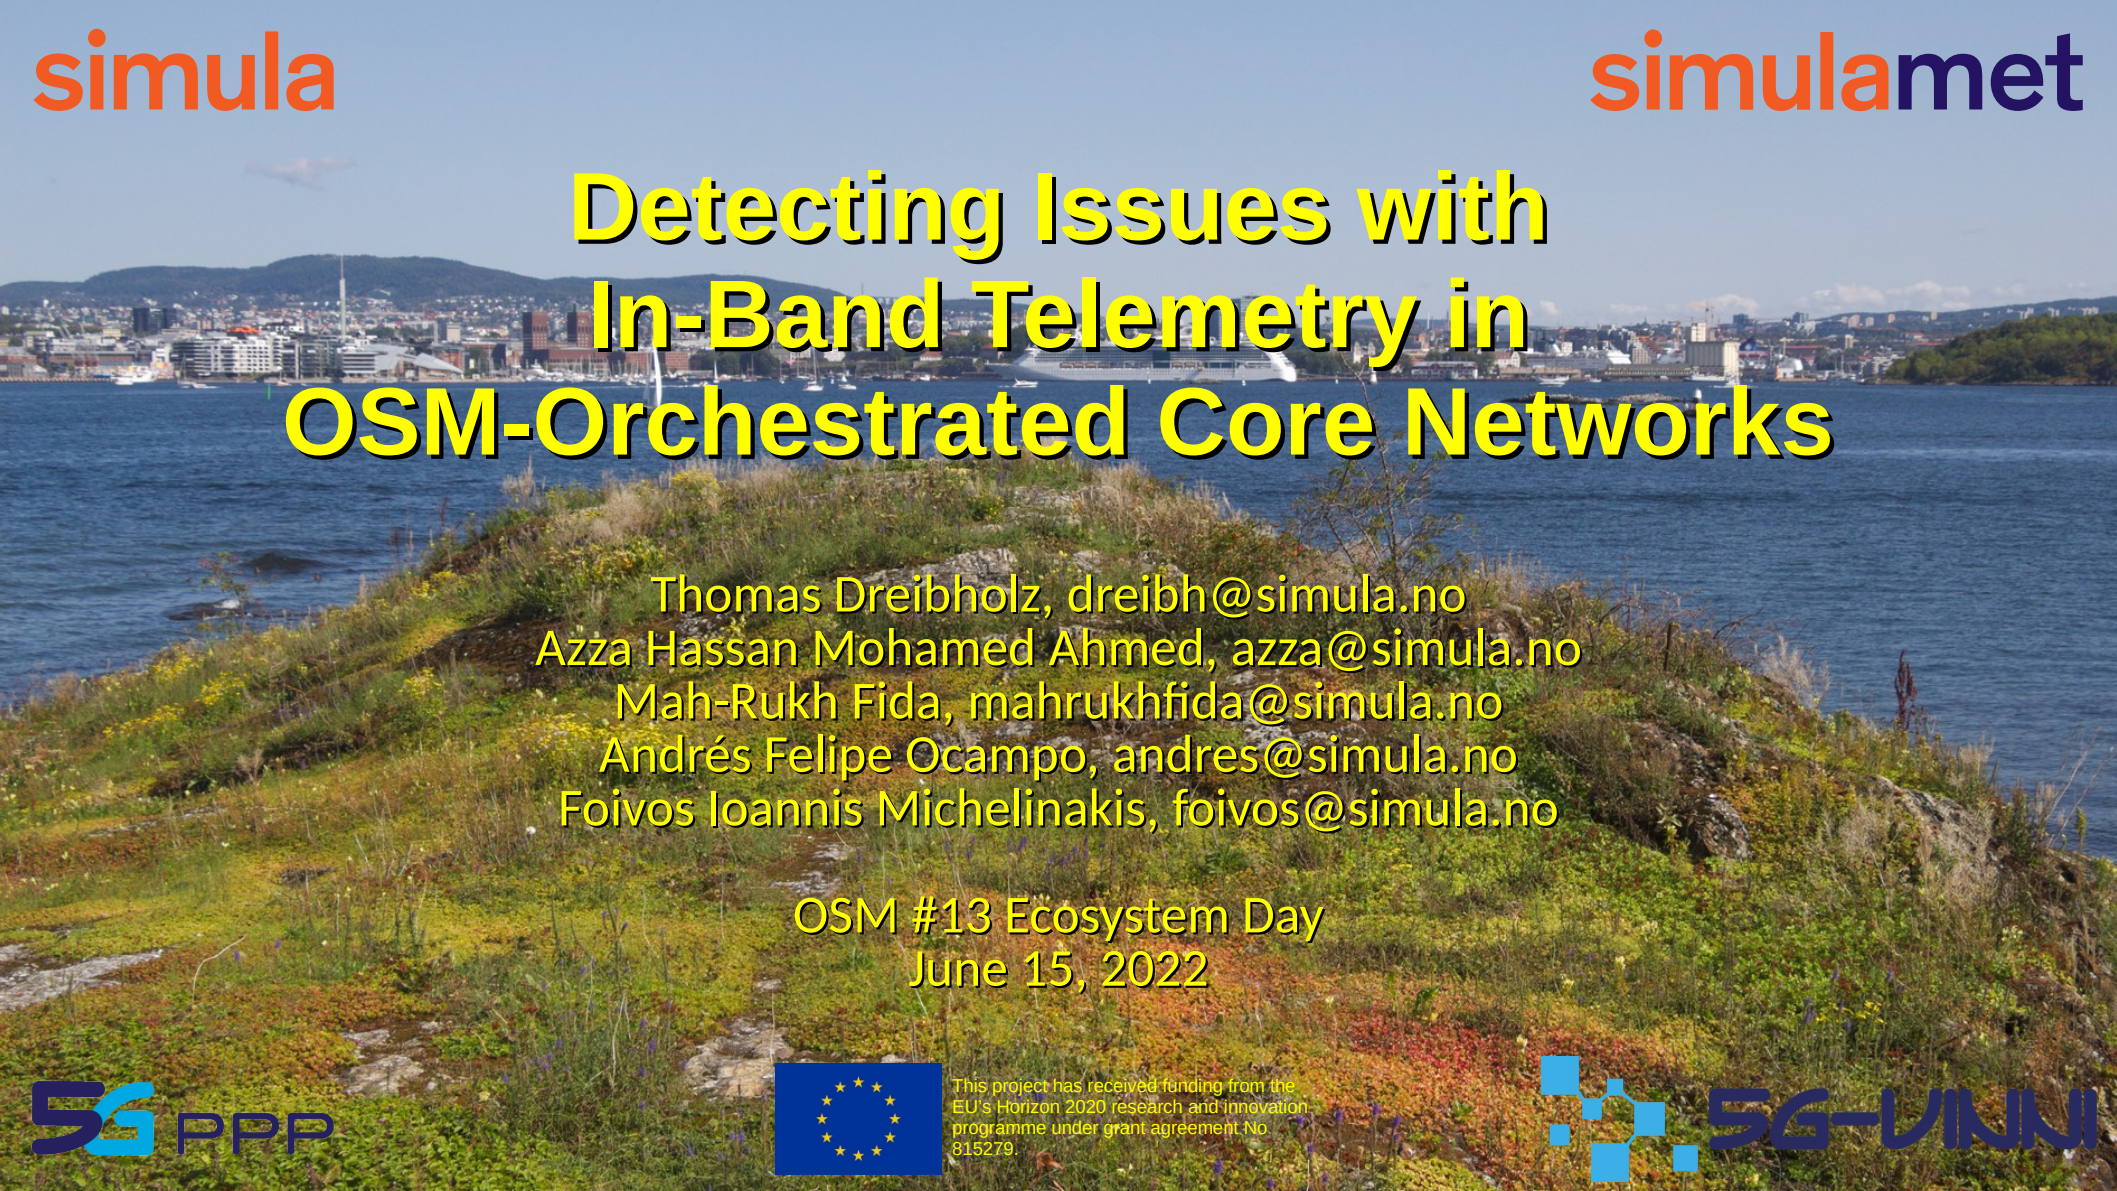

simula

simulamet

Detecting Issues with In-Band Telemetry in OSM-Orchestrated Core Networks

Thomas Dreibholz, dreibh@simula.no
Azza Hassan Mohamed Ahmed, azza@simula.no
Mah-Rukh Fida, mahrukhfida@simula.no
Andrés Felipe Ocampo, andres@simula.no
Foivos Ioannis Michelinakis, foivos@simula.no

OSM #13 Ecosystem Day
June 15, 2022

This project has received funding from the EU's Horizon 2020 research and innovation programme under grant agreement No. 815279.

5G PPP

5G-VINNI

Table of Contents

Motivation

- 4G/5G networks need monitoring:
  Detect problems, evaluate performance, etc.
- Basic performance: metrics in OSM, via Juju and NFVI
  Useful, but not very fine-granular
- In-Band Telemetry:
  - Fine-granular monitoring at packet level
  - Possibility to add custom information into packets
    (additional headers, etc.) and process them elsewhere
- Our approach:
  - Usage of Programming Protocol-independent Packet Processors (P4)
  - Custom P4 software switches as VDUs in our testbed

OpenAirInterface (OAI)

- OpenAirInterface (OAI):
  - Open Source software for EPC and eNodeB (i.e. packet core and base stations)
  - Details: https://www.openairinterface.org
  - 4G LTE available, 5G under development
  - Ongoing work, with many different Git branches
- Idea:
  - Manage OAI setups in OSM (at least, the EPC part)
    - Automatic setup and deployment
  - Easy to add additional features (e.g. Mobile Edge Computing components)
  - Open Source, of course! → https://github.com/simula/5gvinni-oai-ns

Our Testbed Setup

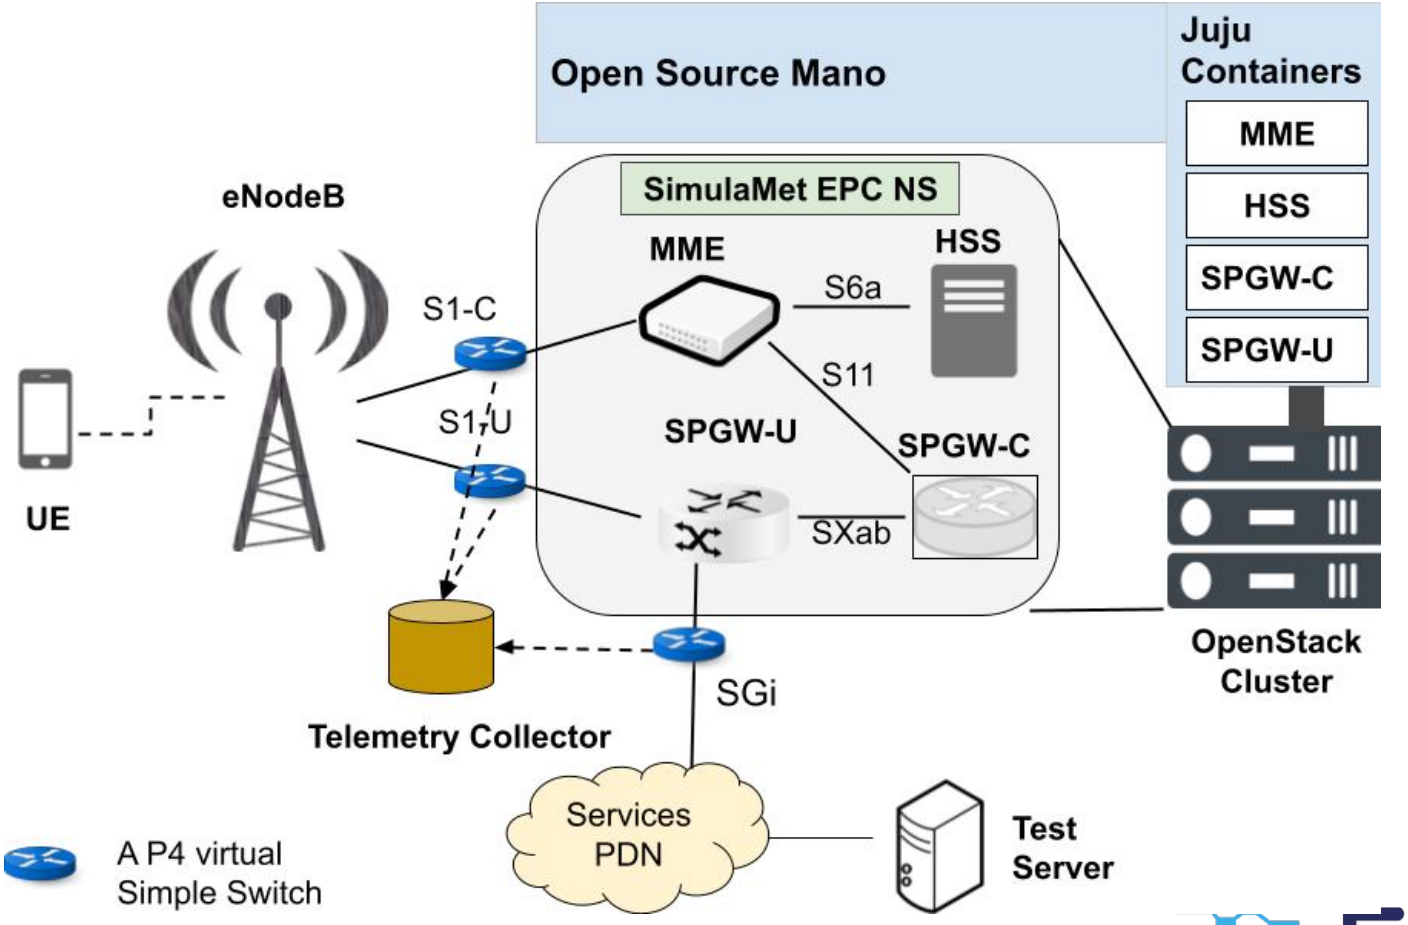

Telemetry at S1-U Switch: Added IPv4 Option Header

```
//Inband Telemetry with IP OPTIONS.................
hdr.ipv4_option.copyFlag = 1;
hdr.ipv4_option.optClass = 0;
hdr.ipv4_option.option= 31;//1 byte
hdr.ipv4_option.swid = 1; //3 bits
hdr.ipv4_option.flowid = meta.flowid;
hdr.ipv4_option.flow_packet_count = (bit<16>) meta.min_count;//2 bytes
hdr.ipv4_option.packets_in_queue = (bit<8>) var_deq_qdepth;
hdr.ipv4_option.hitter = meta.marker; //1 bit
hdr.ipv4_option.queue_timedelta = (bit<26>) var_deq_timedelta;
hdr.ipv4_option.packet_length= (bit<18>)var_packet_length;
hdr.ipv4_option.optionLength =  24; //total number of byte
```

Telemetry at S1-U Switch: P4 Program Example

```
control MyDeparser(packet_out packet, in headers hdr) {
    apply {
                packet.emit(hdr.ethernet);
                packet.emit(hdr.ipv4_outer);//modify in the mirrored packet
                packet.emit(hdr.gre); // add this in mirrored packet
                packet.emit(hdr.udp_outer);//skip this in mirrored packet
                packet.emit(hdr.gtp);//skip this in mirrored packet
                packet.emit(hdr.ipv4_inner); //modify in the mirrored packet
                packet.emit(hdr.ipv4_option);//add this in the mirrored packet
                packet.emit(hdr.udp_inner);
                packet.emit(hdr.tcp_inner);
                }
}
```

Logging of Telemetry by the Telemetry Collector

```
timestamp;flowid;swid;flow_packet_count;packets_in_queue;queue_timedelta;hitter;packet_length
1648111655;154197811627097217071957714537792L;1;4;0;105;0;98
1648111656;240168257070742086464890496 2L;1;5;0;98;0;98
1648111656;240168257070742086463193134 9L;2;5;25;34379;1;1342
1648111658;240168257070742086464890496 2L;1;7;0;225;0;98
1648111658;240168257070742086463193134 9L;2;7;0;170;0;62
1648111658;154197811627096488060830414 93344L;2;7;0;25;0;62
1648111662;154197811627097217071957714537792L;1;9;0;2941;0;98
1648111674;240168257070742086464890496 2L;1;24;0;132;0;98
1648111678;240168257070742086463193134 9L;2;29;0;129;1;62
```

Shown in real time on the collector

Has P4-switch a potential to detect anomaly in end-to-end communication?

Scenario: End-to-End Performance Degradation

Scenario	Description	UE (Performance Tracking Metric)	P4 Switch (S1-U) Telemetry metric
Baseline	TCP downlink data from the test server to UE (using Iperf)	RTT to the test server Jitter to the test Server Transfer (MB) Bitrate (Mbps) Packet loss	dequeue time delta Packets in queue Hitter Packet loss

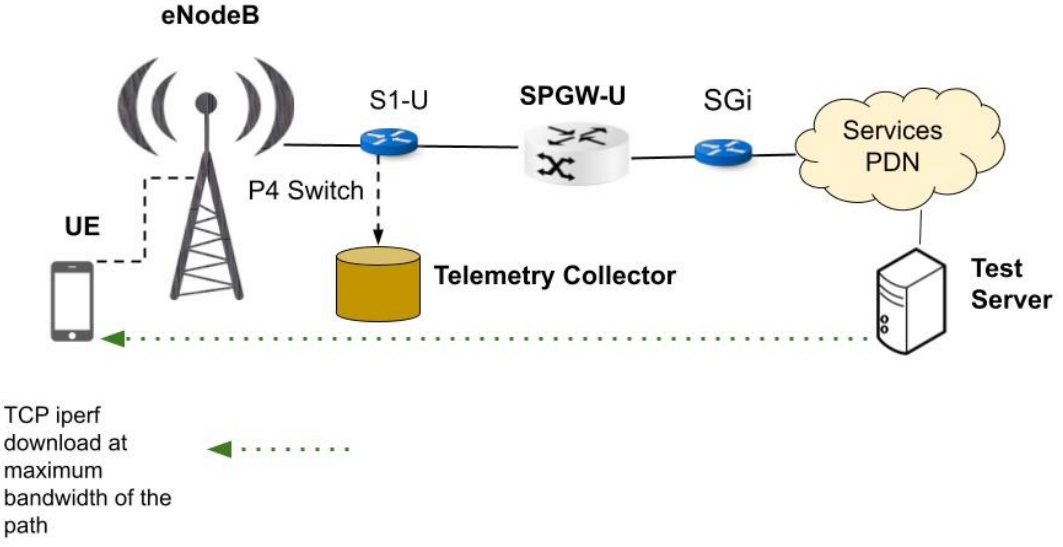

Scenario: End-to-End Performance Degradation

Scenario	Description	UE (Performance Tracking Metric)	P4 Switch (S1-U) Telemetry metric
E2E Anomaly: Bottleneck at S1-U	TCP downlink data from the test server to UE (using Iperf) + TCP downlink data from the SPGW-U to eNodeB (using Iperf)	RTT to the test server Jitter to the test Server Transfer (MB) Bitrate (Mbps) Packet loss	dequeue time delta Packets in queue Hitter Packet loss

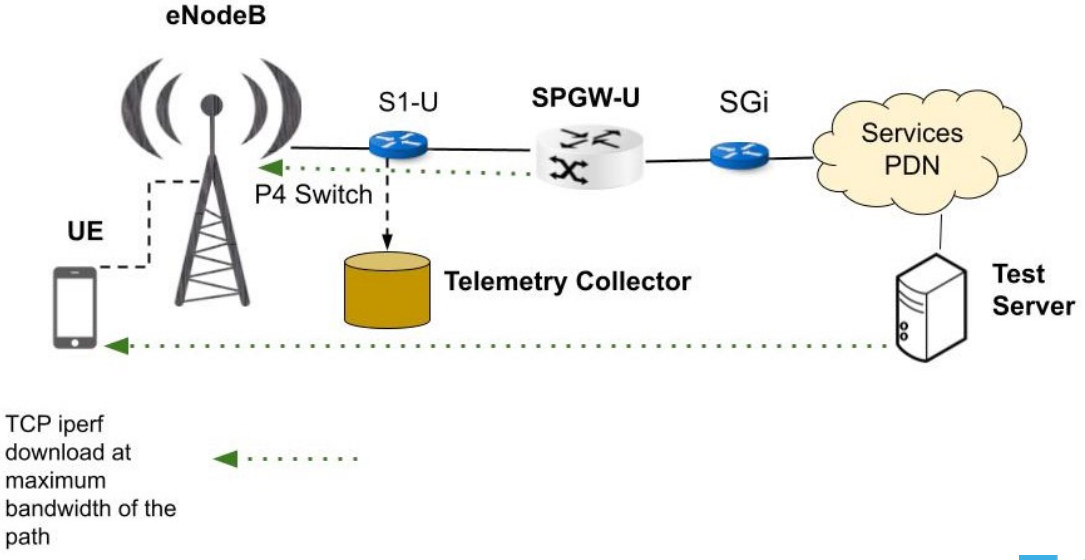

Potential of In-band Telemetry

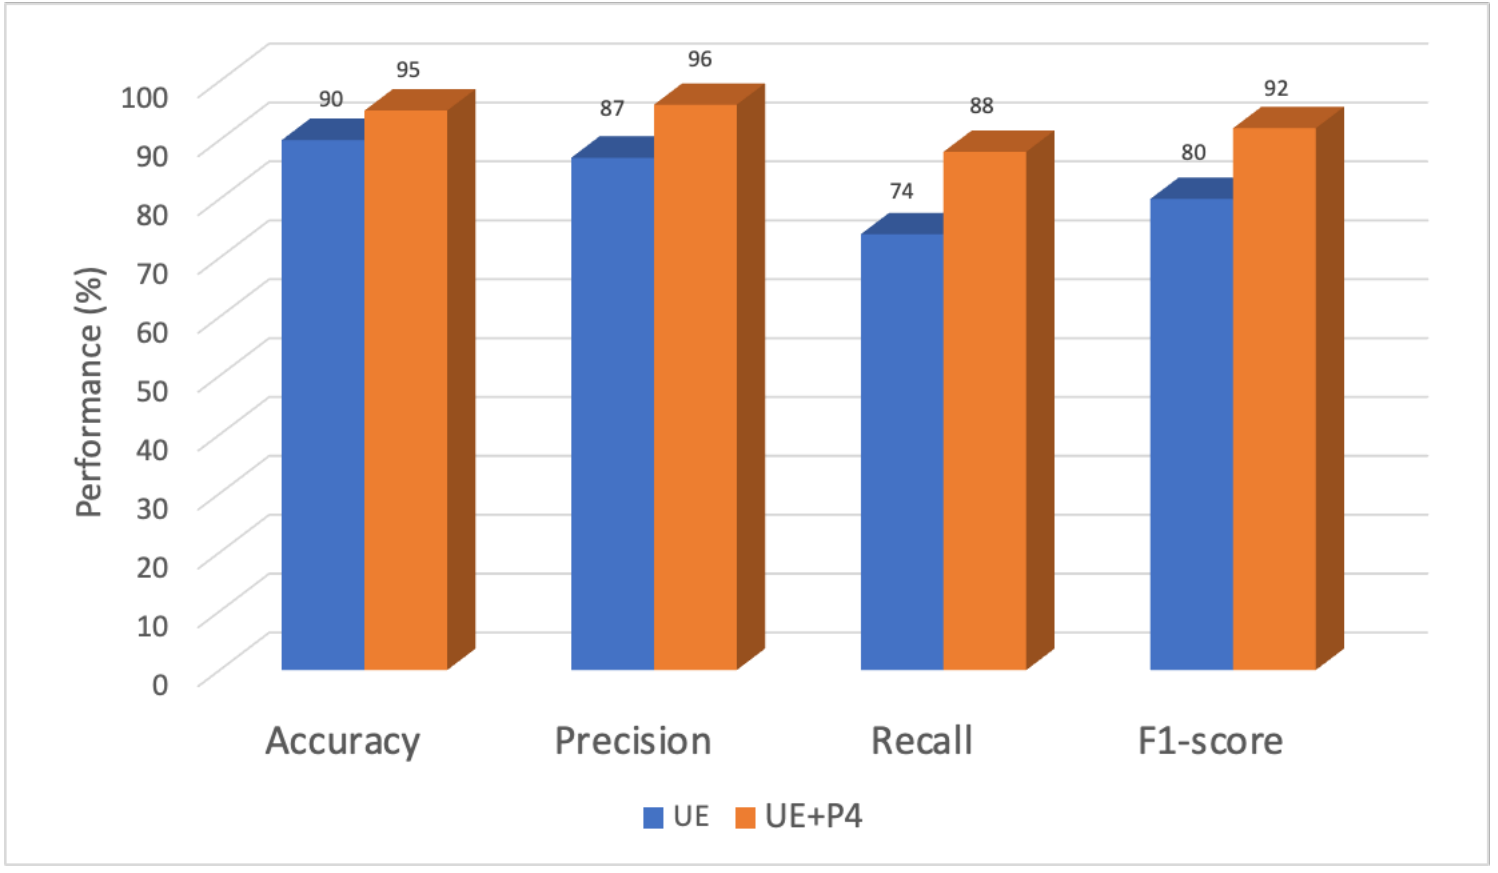

Live Demo!

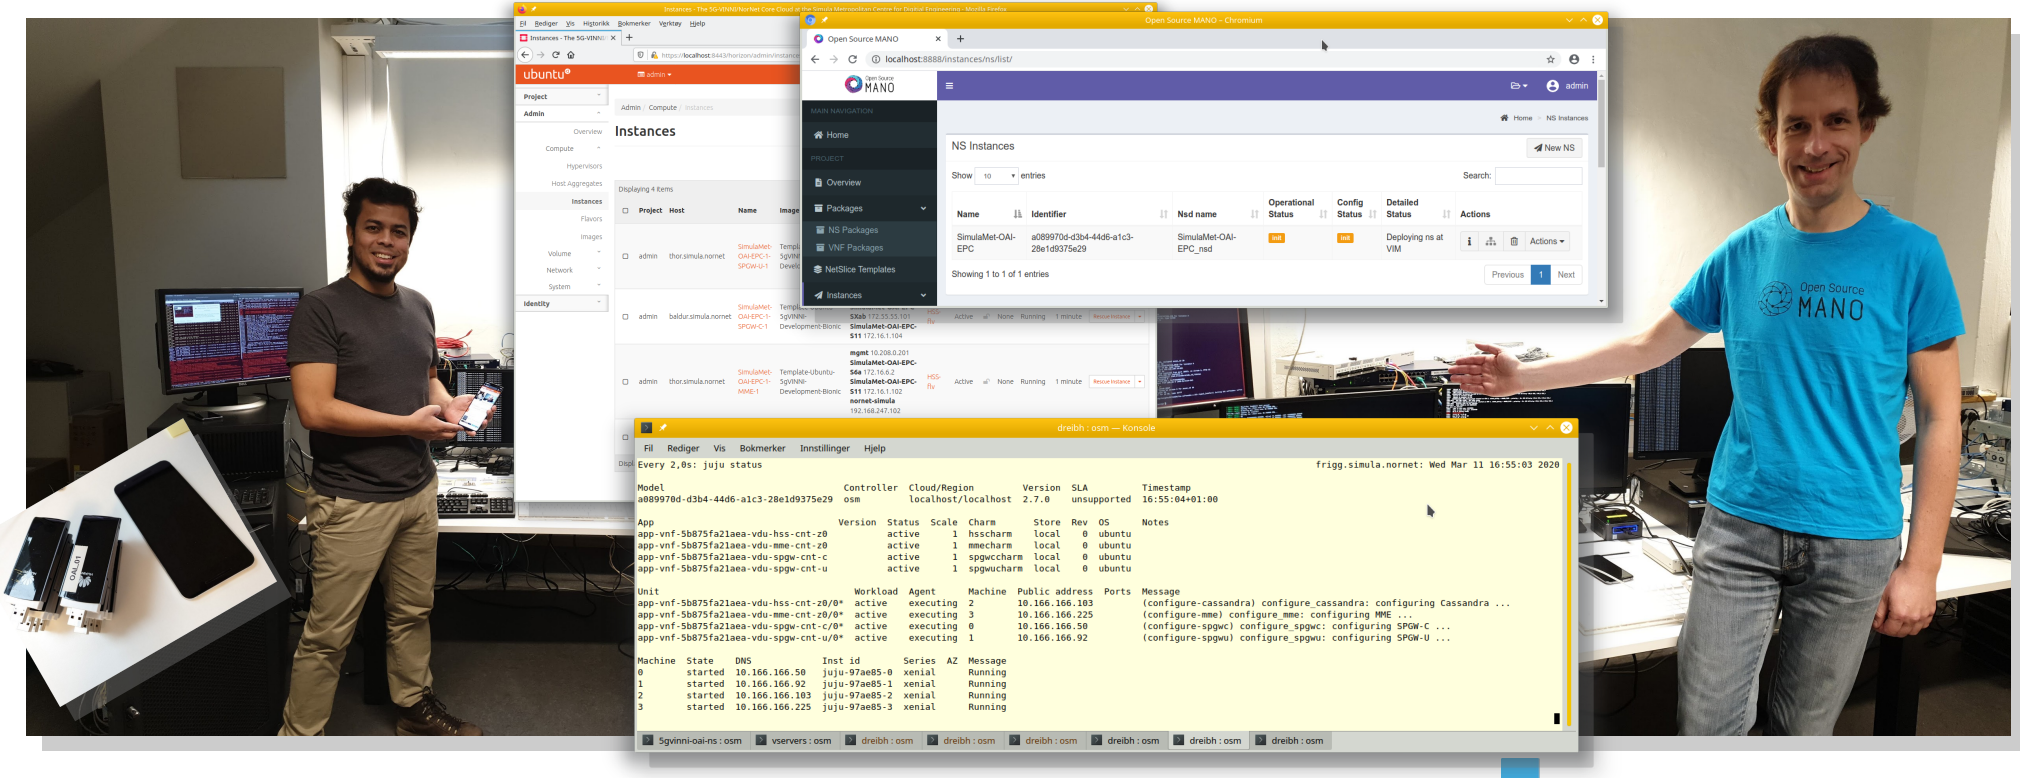
Open Source MANO
NS Instances
Every 2,0s: juju status
Open Source MANO

Any Questions?

Thomas Dreibholz, dreibh@simula.no, https://www.simula.no/people/dreibh

Azza Hassan Mohamed Ahmed, azza@simula.no, https://www.simula.no/people/azza

Mah-Rukh Fida, mahrukhfida@simula.no, https://www.simula.no/people/mahrukhfida

Andrés Felipe Ocampo, andres@simula.no, https://www.simula.no/people/andres

Foivos Ioannis Michelinakis, foivos@simula.no, https://www.simula.no/people/foivos

This project has received funding from the EU's Horizon 2020 research and innovation programme under grant agreement No 815279.