SUPERFLUIDITY

Superfluidity:
A Superfluid, Cloud-Native, Converged Edge System

Call: H2020-ICT-2014-2

Topic: ICT 14 – 2014: Advanced 5G Network Infrastructure for the Future Internet

Project Coordinator: Nicola Blefari Melazzi

Pedro A. Aranda Gutiérrez

pedroa.aranda@telefonica.com

Challenges around the (5G) corner

- ITU-T is planning in the long term beyond LTE [IMT-2020]
- Trends in communications call for
  - No experienced latency
  - Increased Experienced Bandwidth (x100 times more than today)
  - Much connection higher density (O(106) connections/km2)
  - Resulting in increased traffic volume density (O(10Tbps/mk2)
  - Support for much faster mobility (what can I say in the land of the Shinkansen?)
- Looks like we are challenging some laws of nature
  - Spectrum is not infinite => higher spectrum efficiency
  - Much more bits moving around => lower cost per bit
  - Performance should not be bought at the expense of higher energy consumption

Overall Idea

- Run network processing virtualized, on-demand on third-party infrastructure located throughout the network
  - At the core in data-centers
  - At micro data-centers at PoPs in telecom networks
  - At the edge, in RANs next to base stations and at aggregation sites
- Develop technologies to allow such services to be “superfluid”:
  - Fast instantiation times (in milliseconds)
  - Fast migration (in hundreds of milliseconds or less)
  - High consolidation (running thousands on a single server)
  - High throughput (10Gb/s and higher)

Superfluid Architecture

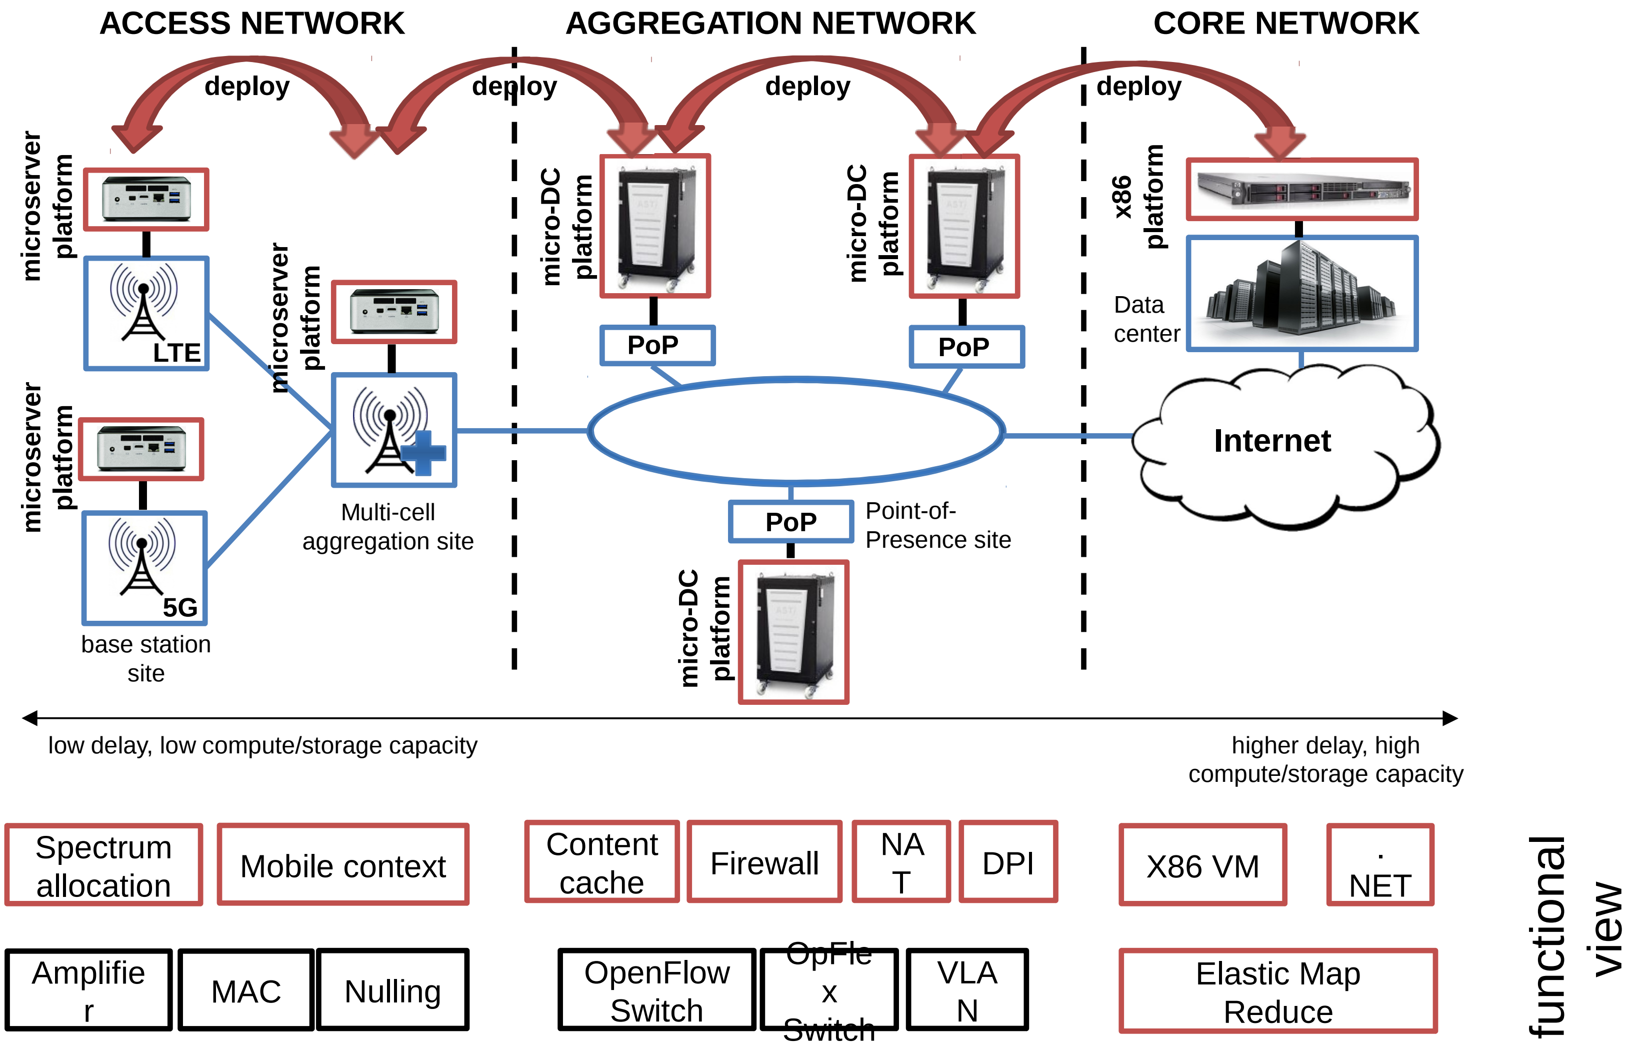

Main Project Goals

- **Converged architecture**: the superfluid platform will abstract the heterogeneity of (1) the underlying hardware and (2) the underlying access technologies
- **Security by design**, to automatically verify that deploying a particular virtualized service won't negatively affect the network or other services
- **Next generation virtualization**: very low instantiation/migration delays, high I/O bandwidth, tiny memory footprints for massive deployments.
- **Heterogeneous hardware acceleration**: leveraging commodity hardware such as FPGAs, GPUs, TCAMs, SoCs, etc), includes dynamic resource allocation algorithms

Why are we significant for the SDN-RG?

- **The Four ‘I’s**
  - **Location independence**: services can be deployed (and relocated) at various networks depending on application needs
  - **Time independence**: fast deployment and relocation in tiny timescales to guarantee service continuity
  - **Scale independence**: transparently scale services in a cloud-like manner, provide massive consolidation
  - **Hardware independence**: the network services (i.e., software) should run on all platforms, irrespective of the underlying hardware
- This can only be achieved with a ‘better’ SDN
  - SDN to the (mobile) edge:
    - Currently SDN is very IP centric
    - To cope with the requirements of the 5G mobile edge, SDN has to leave that confort zone

Sample Use Cases

- Next generation emergency services
- Minimum-delay cloud storage
- Localized services (e.g., gaming, video conferencing, etc.)
- Edge offloading (e.g., ad blocking, firewalling, etc.)
- On-the-fly monitoring
- DDoS Filtering
- Virtual CDNs
- Virtual CPEs

Who and where are we?

Partners

Consorzio Nazionale Interuniversitario per le Telecomunicazioni	IT
Alcatel Lucent Bell Labs France	FR
Alcatel Lucent Israel	IL
British Telecom	UK
Citrix	GR
EBlink	FR
Intel Ireland	IE
NEC Europe	UK
OnApp	UK
Portugal Telecom Innovation and Systems	PT
Red Hat	IL
Telcaria	ES
Telefonica I+D	ES
Unified Streaming	NL
University Ben Gurion	IL
University of Liège	BE
University of Technology Dresden	DE
University Politehnica of Bucharest	RO

Thank you.
Questions?

References

[IMT-2020] http://www.itu.int/dms_pub/itu-r/oth/0a/06/R0A0600005D0001PDFE.pdf