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

Felipe Huici
felipe.huici@neclab.eu

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

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

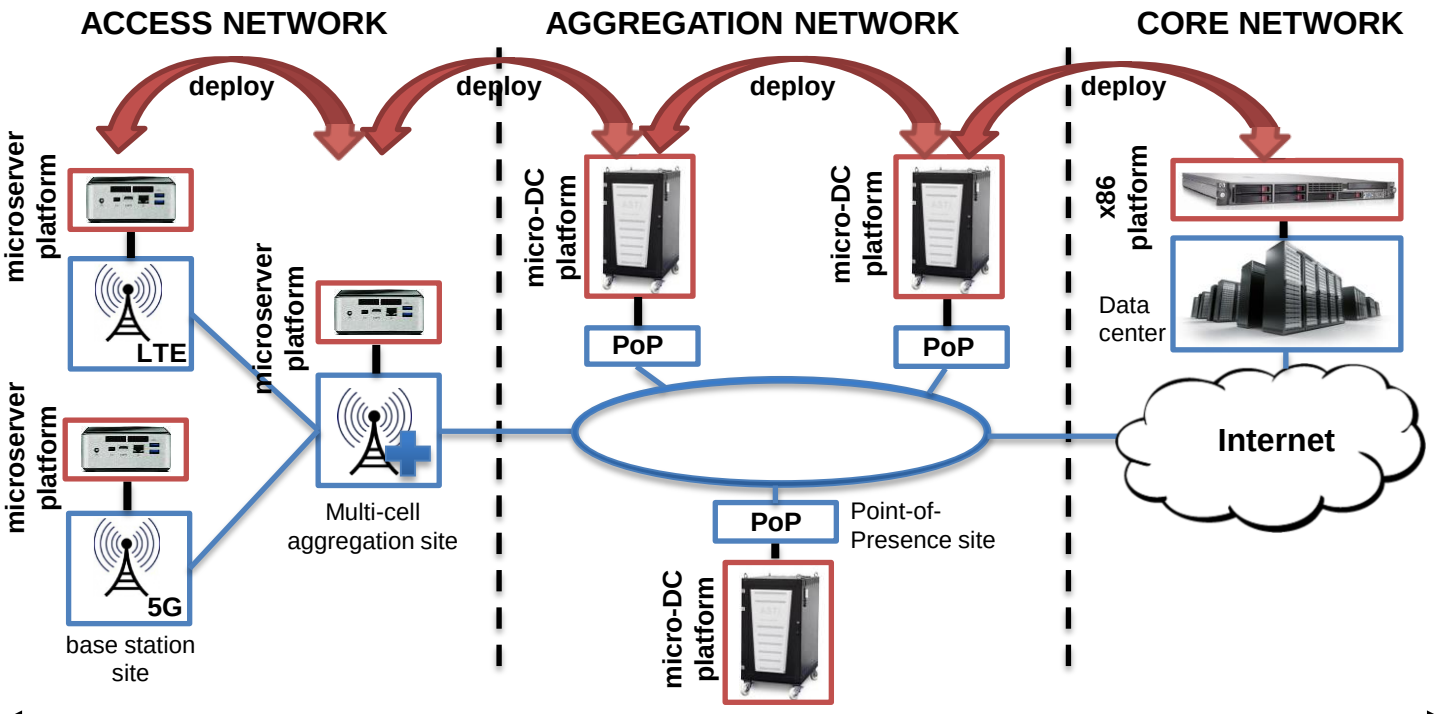

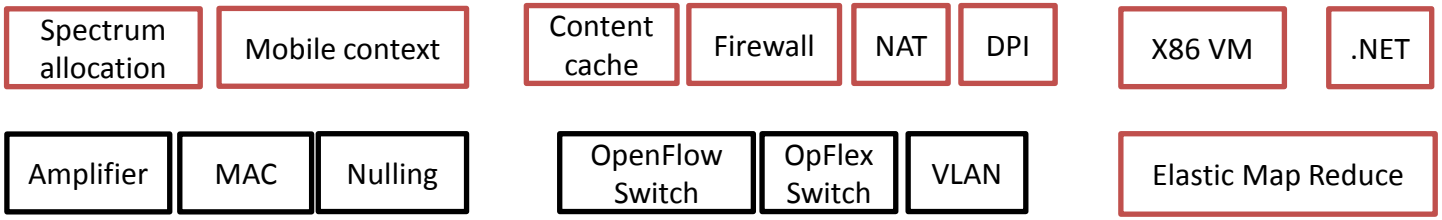

functional view

Target Test Platforms

- x86 servers (AMD, Intel, ARM64 if available)
- Microservers (well suited to lightweight virtualization, edge deployments)
  - Small physical size → Can deploy almost anywhere
  - Low power consumption → Can be run off of a USB battery
  - Passive cooling only → Reduces energy consumption, no noise
  - Low cost → $30-$200

Intel Edison: 61/29/12 mm

AMD Gizmo2: 10.16/10.16 cm

Rasberry Pi B+: 85/5617 mm

Wandboard i.MX6: 95/95 mm

Minnowboard: 99/74 mm

Odroid-U3: 87/52/29 mm

Superfluid Cloud Properties – 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

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)

Sample Use Cases

- 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

A Small Taste…

Towards Superfluidity: Minimalistic Operating Systems

- Observation: general-purpose operating system contain a large amount of functionality that is not always needed (e.g., a virtualized NAT does not need a USB driver)
  - Aim: reduce the overheads inherent in general purpose operating systems
- Expose POSIX-like or part of a POSIX API in order to support “legacy” applications
- Further optimize through *specialization*: build only one or a few applications on top of the operating system
  - Often a single memory space: *no* kernel/user-level divide

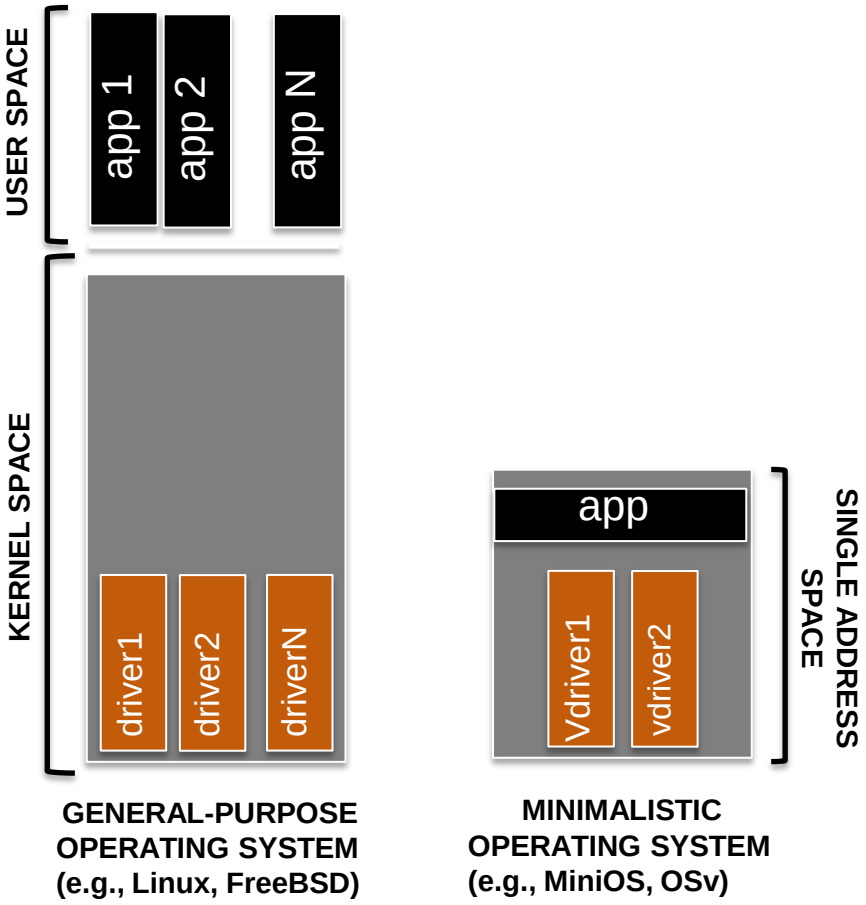

Towards Superfluidity (on Xen)

- Improvements to the Xen toolstack for faster domain creation
  - Without these, creating, e.g., 4,000 VMs takes 5 days
- Improved console backend polling mechanisms
- Re-implemented Xenstore core service
  - Improved boot times by 40%

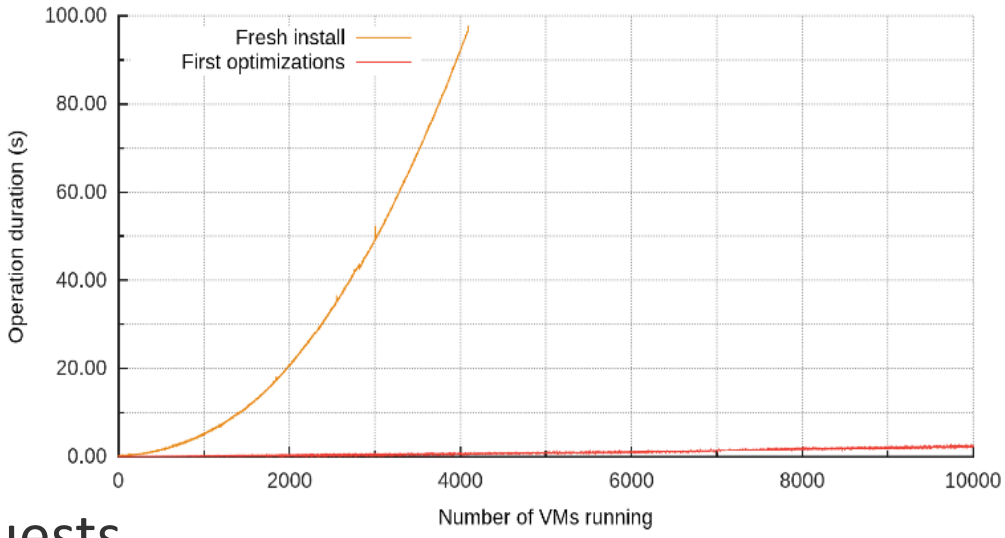

- Created new Xen toolstack for para-virtualized, specialized guests
  - Enables small footprint guests to boot in only 10 milliseconds!
- Summary: can boot > 10K VMs on a single x86 server

Thank you. Questions?