SUPERFLUIDITY

A SUPER-FLUID, CLOUD-NATIVE, CONVERGED EDGE SYSTEM

Research and Innovation Action GA 671566

DELIVERABLE D8.8:
FINAL REPORT ON INNOVATION AND EXPLOITATION ACTIONS

Deliverable Type:	Report
Dissemination Level:	Public
Contractual Date of Delivery to the EU:	31/03/2018
Actual Date of Delivery to the EU:	08/05/2018
Workpackage Contributing to the Deliverable:	WP8

Editor(s):	George Tsolis (CITRIX)
Author(s):	Bessem Sayadi (NOKIA FR), Francisco Fontes (ALB), Diego Lopez (TID), Michael J. McGrath (INTEL), Luis Tomas Bolivar (REDHAT), Elisa Rojas (TELCA), Claudio Pisa (CNIT), John Thomson (ONAPP), Erez Biton (NOKIA IL), Philip Eardley (BT)
Internal Reviewer(s)	Stefano Salsano (CNIT)

Abstract: This deliverable is the final report on the innovation and exploitation actions, on the overall progress achieved and performance against the relevant plan, as well as on the awareness of the innovation and exploitation results that have been achieved overall by the project. Moreover, it describes the impact that these actions have delivered, as well as further innovation that the project will be driving after its completion. This report supersedes deliverable D8.5.

Keyword List: Innovation, Exploitation

INDEX

List of Tables

Glossary

TERM	DEFINITION
Innovation	The act or process of introducing new ideas, technologies, or methods.
Exploitation	Direct or indirect utilisation of foreground in further research activities other than those covered by the project, or for developing, creating and marketing a product or process, or for creating and providing a service, or for using them in standardisation activities. (Article 28.1 of GA)
Results	Any tangible or intangible output of the Action, such as data, knowledge and information whatever their form or nature, whether or not they can be protected, which are generated in the Action as well as any rights attached to them, including Intellectual Property Rights.
General Public	Audience composed by people, groups and organizations which do not have specific competences/knowledge in the matter of the project but only a general interest at the application of results in society at a large.

Table 1: Glossary.

1 Introduction

1.1 Deliverable Rationale

This deliverable provides the final report on the activities of Task 8.3 (T8.3) of Work Package 8 (WP8), which specifies:

Work Package 8: Communication, Dissemination, Standardization and Exploitation

Task 8.3: Innovation and Commercial Exploitation

This task will be in charge of Innovation and Commercial Exploitation management. It will include the following activities:

- Overall project responsibility for identifying and driving the innovation and commercial potential for the technical work and goals of the project.
- Ensure of a common approach for Innovation and Commercial Exploitation in all work packages while maintaining the overall vision of the Project.
- Identification of potential innovation and commercial exploitation inhibitors in and between the work packages during the course of the technical activities.
- Preparation of proposals for the Management Board on innovation concepts.
- Preparation of summaries for the periodic reports and annual reports, and for the reviews.
- Delivering of relevant presentations both internally and externally to the industry.

1.2 Executive summary

1.2.1 Deliverable description

This deliverable is the final report on the innovation and exploitation actions taken, on the overall progress achieved and performance against the relevant plan (Deliverable D8.3), as well as on the awareness of the innovation and exploitation results that have been achieved overall by the project.

Moreover, it describes the impact that these actions have delivered, as well as further innovation that the project is expected to be driving after its completion. In summary, this report will include the measurable exploitation results that have been achieved and their stimulus towards further growth.

This report supersedes Deliverable D8.5.

1.2.2 Deliverable structure

Section 2 covers Innovation Management. More specifically, Section 2.1 provides a summary of the achievements of the project until M33. Section 2.2 outlines the innovation pathways the project has actively pursued. Section 2.3 provides a final update on the progress beyond the state of the art. Last but not least, Section 2.4 summarises the innovation potential and anticipated impact of the project.

Section 3 covers Exploitation Activities and Plan. Section 3.1 outlines the updated exploitation plans of the individual partners, while Section 3.2 summarises the exploitation plan of the consortium as a whole. Section 3.3 lists the Proof of Concept (PoC) activities planned by the project, as well as their status.

Finally, Section 4 provides a quick summary of measurable project innovation and exploitation results. The deliverable concludes with a list of References, related to both the initial state of the art (original research bibliography) and the published project research results that advanced the state of the art (project publications are highlighted in bold text).

2 Innovation Management

Activities in this area aim to fulfil the Scientific and Technological Vision of the SUPERFLUIDITY project (please refer to section 3.2 of deliverable *D1.2 Project Vision and Roadmap*).

To foster innovation, the project must first implement the appropriate management structure and provisions for frictionless collaboration between the partners. Sections 2.1.1 and 2.1.2 of deliverable *D8.3 Innovation and Exploitation Plan (Updated Version)* provided more information on this matter. The measures we have identified for counteracting the potential inhibitors (see Section 2.1.3 of deliverable above) are in place and play an important role towards maximizing the potential impact.

Secondly, the project has to keep pace with the ongoing scientific research relevant to the technological objectives. Section 2.2 of deliverable D8.3 outlined the state of the art across the various research domains of interest, which the project continues to review and closely monitor.

Thirdly, the project has to advance the state of the art, in research domains that are aligned with the project scientific and technological vision. In the form of a summary, Section 2.1 describes the achievements and impact of the project so far. Section 2.2 outlines the innovation pathways the projects is taking. Section 2.3 then reports on the progress towards advancing the state of the art.

Finally, and as an extension to the research directions and innovation pathways covered above, Section 2.4 summarises the innovation potential of the project. Most importantly, this is associated to specific market segments of interest to the European Communications Services Providers (CSPs), industrial partners and SMEs of the project, where the innovations of SUPERFLUIDITY are expected to generate more impact, resulting in competitive advantages.

The wealth of project activities reported in deliverables *D8.6 Final Report on Communication & Dissemination Actions* and *D8.7 Final Report on Standardisation & Open Source Contributions*, which we will not replicate herein, should also be considered a measure of innovation and a driver of exploitation of the project.

2.1 Project Achievements

The SUPERFLUIDITY project is focused on the transformation of Telecommunications and Service Providers infrastructures to Cloud Native infrastructures in order to bridge the Cloud and Telco worlds. The vision of the project consists of a "superfluid" 5G network infrastructure, characterized by four key properties: 1) location-independence where services can be deployed (and relocated) on various networks depending on application needs, 2) time-independence where deployment and migration of services is near instantaneous, 3) scale-independence where services scale transparently in a cloud-like manner, providing massive consolidation, 4) hardware-independence where services are developed and deployed with high performance irrespective of the underlying hardware.

Following this vision, the SUPERFLUIDITY project worked on a set of key innovations. Firstly, network services are decomposed into Reusable Function Blocks (RFBs). Decomposing functions into RFBs facilitates the re-use of components. The project considers heterogeneous platforms for the deployment and execution of RFBs, called RFB Execution Environments (REE). The composition and orchestration of RFBs on the REEs is modelled using the RFB Description and Composition Languages (RDCL). RFBs and REEs respectively represent a generalization of the concepts of VNFs (Virtual Network Function) and NFVI (NFV Infrastructure) defined in current ETSI NFV architecture, while the RDCLs are the generalization of the descriptors of Network Services and VNFs. With respect to the state of the art, the introduction of RFBs, REEs and RDCLs proposed by SUPERFLUIDITY provides a higher level of abstraction, including the support of heterogeneous execution platforms and the possibility to decompose functions with finer granularity (i.e. into smaller components). In particular, the decomposition can be iterated as RFBs can be decomposed in RFBs resulting in a "nested" composition model.

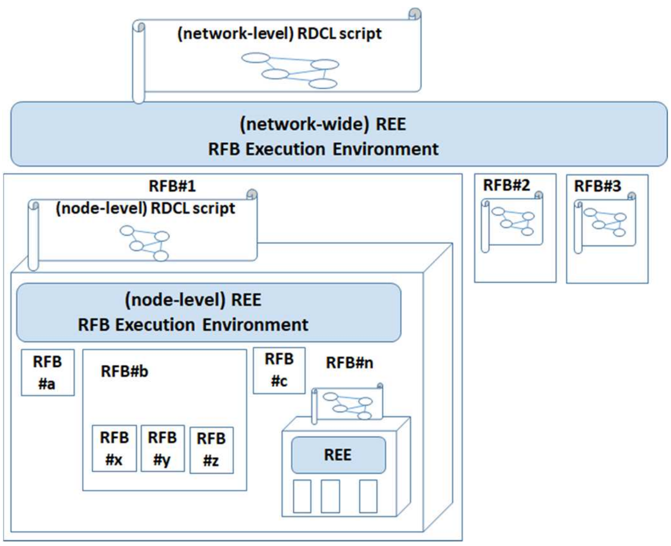

Figure 1: RFB, REE and RDCL abstractions proposed by SUPERFLUIDITY

A key benefit of using RFBs is that by decomposing complex functions into their constituent parts it is possible to re-utilize their atomic functionalities to quickly compose new services. More information can be found in the SUPERFLUIDITY D3.1 - Final system architecture, programming interfaces and security framework specification.

The project has investigated Micro-VNFs, small and highly specialised VNFs that can be supported by the Unikernel virtualization technology. Unikernels offer very good performance in terms of memory footprint and instantiation time and have very good isolation and security properties (better than containers). In particular, the project investigated ClickOS, a Xen-based Unikernel tailored for NFV appliances, able to provide highly efficient raw packet processing. In our tests, Unikernel instances have been demonstrated to have a small footprint (around 5 MB of memory when running), an instantiation time of a few milliseconds, capable of processing up to 10Gb/s of traffic and do not require a persistent disk drive to work. In addition, they benefit from the isolation provided by the Xen Hypervisor and the flexibility offered by the Click modular router.

The project also investigated the combination of virtualisation technologies (VMs and containers) in the same service infrastructure and tackled a set of challenging problems related to the manner in which they will be orchestrated and networked together. In this area, the project contributed to open source frameworks such as ManageIQ for service deployment and to Kuryr to provide the networking capabilities necessary to support the deployments of VM's and containers together.

The SUPERFLUIDITY project has integrated the MEC (Mobile Edge Computing) concept into its architecture, designing and implementing a modular MEC prototype using SDN/NFV technologies. The MEC is integrated in an EDGE cloud deployment with the first prototype of a Cloud RAN that includes a front-haul re-programmable via SDN, and interacting with a decomposed core network. The different Cloud RAN and CORE components are deployed as Docker containers.

The SUPERFLUIDITY vision of a highly dynamic deployment of RFBs to support 5G services requires the appropriate use of telemetry, analytics and data visualisation. Telemetry or monitoring helps to provide useful insights into platform issues, which can significantly influence the performance of user-provided service-level KPIs. In fact, the type, quantity, quality and configuration of the exposed metrics by virtualized functions or host cloud environments can be limited leading to significant challenges. Metrics are typically tied to a specific type of virtualisation implementation approach such as virtual machines (VMs) and can have scalability challenges. In the context of SUPERFLUIDITY, a scalable and flexible telemetry platform is required to support the diverse virtualised environments i.e. VMs/containers and bare metal (non-virtualised environments). To address this challenge, the project has adopted standalone telemetry agents to provide a wide range of metrics across different virtualisation methods and for different use cases i.e. operation service monitoring and service characterisation.

The characterisation methodology relies on a structured experimental approach focusing on specific service characteristics called TALE (Through, Anomalies, Latency, Entropy), which leverages a full stack monitoring approach for the collection of metrics. Collected metrics relate to both the physical and virtualised compute/storage/network environments and the actual service under test in an operational context. The snap framework has been adopted by the project due to its set of capabilities, which address many of the requirements identified by SUPERFLUIDITY 5G environments. For example the snap telemetry platform supports dynamic reconfiguration without interruption, can be used in "tribes" for automated replication of agent configurations i.e. change on one, replicate to many, which is an important feature for scalability.

Another critical aspect targeted by SUPERFLUIDITY is related to the formalised validation of deployments. In this area, SUPERFLUIDITY considers a symbolic execution tool for dataplane verification called SymNet and is working on:

- Languages that allows operators to express their policies easily and with a compact representation.
- A verification tool that performs symbolic execution guided by the operator policy, in order to reduce the number of explored paths.
- Provably correct transformations from the SEFL language (used by SymNet) to dataplane languages such as P4, ensuring that the verification results are accurate.

The SUPERFLUIDITY project has demonstrated its main achievements in an integrated testbed. The final demonstration has been structured in the form of a set of scenes, meant to demonstrate the key features claimed by the project. The demonstrator is built using two integrated testbeds (connected using a VPN), which are geographically dispersed and host different elements of the end-to-end demonstrator. Testbed #1 is located at Nokia France premises and provides a hardware and wireless platform supporting the demonstration of a number of the key innovations (e.g. Cloud RAN, RFB decomposition …) developed by the project. It acts as the EDGE cloud. Testbed #2 is a hardware platform located at BT's UK premises, consisting of 5 servers and a switch, which supports flexible virtualisation experiments and demonstrations. It acts as Central Cloud.

2.2 Innovation Pathways

Table 2 outlines the research areas targeted by SUPERFLUIDITY, as well as the objectives of the project for each of these areas. Line items in *italic* text are additions to the initial scope of the project.

Research Area		SUPERFLUIDITY Goals
Cloud Networking	NFV Virtualised Environment	- Achieve 10-40 Gbps throughput with virtualised and software based packet processing. - Achieve tens of ms. (or lower) service instantiation. - Achieve massive consolidation and migration of services.
	Software Radio and DSP Virtualization	- Unify and abstract the protocol stack and the hardware platform to instantiate multiple vendor protocols. - Open radio platform and enable third parties to design the protocol stack. - Design portable runtime dataflow engines, enabling simultaneous seamless execution of multiple protocols.
	High Perf. Software-Based Packet Processing	- Achieve high performance with commodity hardware. - Increase the level of flexibility by going beyond the paradigms of proprietary, embedded systems. - Enable agile network function deployment.
	Automated KPI to Platform Metrics Mapping	- *Develop an automated methodology to identify platform features which most significantly influence user provided KPI's for a given workload* - *Reduce the number of metrics which must be captured by an order of magnitude to monitor a KPI* - *Develop an approach/model which can predict user provided service KPI's*
Network Service Decomposition and Programmability		- Introduce program abstractions specifically targeted to 5G functions. - Combine block-based composition abstractions (such as those exploited in Click routers [KOH2000], or emerging in the ETSI NFV work on service chaining) with event-driven programming paradigms such as basic match/action based approaches or more powerful stateful abstractions based on extended finite state machines.
Cloud RAN and Mobile Edge Computing		- Enable modular "hot" replacing of eNB functions (such as scheduling) and allow migration of such functions between edge clouds and the antenna subsystem, so as to balance algorithmic complexity with front-haul capacity. - Enable the migration of non-RAN functions (like local caching and CDN) to the edge cloud, to maximise their performance.

	- *Decompose video streaming functions to exploit the combined core and edge architecture, reducing the transfer of video data in the back-haul.*
Automated Security and Correctness	- Provide a pre-deployment checking system to ensure that virtualised network services do not negatively affect the network nor other tenants; the system has to be both scalable and stateful, able to model most types of services. - Implement a post-deployment system that will learn the behaviour of traffic and detect any anomalies, thus providing a further security mechanism in cases where the checking system does not have information about the processing performed by a network function, or when static analysis is inaccurate.

Table 2: SUPERFLUIDITY Innovations and Goals

2.3 Progress Beyond the State of the Art

In this section we cover what the project's progress beyond the state of the art was envisioned to be, for each area of Table 2. Also, we report project results that demonstrate outcomes to that direction.

One general area of interest to the project is to evaluate the viability of deploying future 5G networks, and how it might benefit from SUPERFLUIDITY innovations. In ***[CHI2016]****, SUPERFLUIDITY researchers investigate the possibility of deploying 5G networks in rural and low-income zones. After considering the main challenges that need to be faced, we identify the main pillars to follow, in order to deploy 5G networks in such zones, as well as a proposal of a future network architecture. We expand on this further in* ***[CHI2017a]****, which was published in the IEEE Communications Standards Magazine. Our solution investigates the possibility of mounting Remote Radio Heads (RRHs) on top of Unmanned Aerial Vehicles (UAVs), as well as Large Cells (LCs) to increase the coverage range. In addition, 5G-nodes are powered by solar panels and batteries. Preliminary results, obtained over three representative case studies (located in Italy, Cook Islands, and Zimbabwe), show that providing connectivity in rural and low-income areas by means of the proposed 5G architecture is feasible. At the same time, we show that the monthly subscription fee paid by the users can be kept sufficiently low, i.e., less than 1 [EUR/month] in low-income areas, and around 11 [EUR/month] in rural regions.*

As a generalization of the above, in ***[CHI2017b]*** *we target the evaluation of a Superfluid 5G network from an economic point of view. The considered 5G architecture has attributes, such as flexibility, agility, portability and high performance, as advocated by the project. The proposed economic model, allows to compute the CAPEX, the OPEX, the Net Present Value (NPV) and the Internal Rate of Return (IRR). Specifically, we apply our model to estimate the impact for the operator of migrating from a legacy 4G to a 5G network. Our preliminary results, obtained over two realistic case studies located in*

Bologna (Italy) and San Francisco (CA), show that the monthly subscription fee for the subscribers can be kept sufficiently low, i.e., typically around $5 per user, while allowing a profit for the operator.

2.3.1 Cloud Networking

2.3.1.1 Cloud-based, Virtualised Network Services

SUPERFLUIDITY will push the boundaries of the previous work with the goal of developing mechanisms in support for a superfluid, virtualised network architecture. Specific goals include, but are not limited to, achieving high throughput (10-40Gb/s and higher throughput), near-instantaneous virtualised service instantiation and migration (in tens of milliseconds or less) and massive consolidation of services (e.g., thousands of virtual machines on a single x86 server).

Network and function Management and Orchestration (MANO) is of paramount importance for SUPERFLUIDITY, therefore we aim to enhance existing, yet still basic orchestration platforms [MIJ2016b]*. In* [VEN2016] *we deal with the management of the virtual computing resources for the execution of Micro VNFs. This functionality is performed by the Virtual Infrastructure Manager (VIM) in the NFV MANagement and Orchestration (MANO) reference architecture. We discuss the VIM instantiation process and propose a generic reference model, starting from the analysis of two Open Source VIMs, namely OpenStack Nova and Nomad. We implemented a tuned version of the VIMs with the specific goal of reducing the duration of the instantiation process. We realized a performance comparison of the two VIMs, both considering the plain and the tuned versions. The tuned VIMs and the performance evaluation tools that we have employed can be downloaded from our repository. The comparison was thereafter expanded to add OpenVIM, which is part of ETSI Open Source MANO. In* [LUN2018]*, presented at the ETSI Open Source MANO Conference, we describe the proposed extensions to OpenVIM to support ClickOS Unikernels and Xen. On the basis of a scenario we have implemented, that can combine Unikernels and regular VMs in the same Network Service or VNF extending OpenVIM, we describe how we have extended the ETSI NFV models and OpenVIM. In particular, we provide the details of the OpenVIM descriptor extensions to support Unikernels. As a background information, we discuss Unikernels and their orchestration aspects, focusing on ClickOS Unikernels. We have adapted 3 VIMs (OpenStack, Nomad, OpenVIM) to support ClickOS Unikernels and report the performance evaluation of VM instantiation time.*

Recently, with new hardware architectures such as Reconfigurable Match Tables and languages such as P4, the SDN community has started to bring line-rate data plane programmatic flexibility inside switching chipsets. Starting from the original OpenFlow's match/action abstraction, most of the work has so far focused on key improvements in matching flexibility. Conversely, the "action" part, i.e. the set of operations (such as encapsulation or header manipulation) performed on packets after the *forwarding decision, has received way less attention: The OpenFlow community has limited to standardize the set of supported actions, whereas their implementation has been delegated to each specific vendor/device. Our goal is to move beyond the idea of "atomic", pre-implemented, actions, and rather make them programmable while retaining high speed multi-gbps operation. In [PON2017] we propose a domain-specific HW architecture, called Packet Manipulation Processor (PMP), able to efficiently support micro-programs implementing such actions. We describe three non-trivial use cases (tunnelling, NAT, and ARP reply generation), and assess the relevant throughput performance. Based on our analysis, the most demanding actions, in terms of processing power, are those related to the manipulation of the data inside the packets, i.e. the set of operations (such as encapsulation or header manipulation) performed on packets after the forwarding decision. These encapsulation and header rewriting actions are key elements to put together heterogeneous networks and to enable flexible slicing, as envisioned by the forthcoming 5G network architecture. In [BRU2018] we focus on the design of a Very Long Instruction Word (VLIW) processor, based on a custom instruction set, able to perform efficient packet processing operations granting multi-gbps throughput. In particular, we provide details of the Packet Manipulator Processor (PMP) architecture and its I/O interfaces, which have been designed to accomplish packet manipulation tasks, we discuss the throughput analysis of the three abovementioned use cases, and present FPGA synthesis results.*

Another major challenge with Network Functions Virtualisation (NFV) is ensuring protection of Virtual Network Function (VNF) resources against "Noisy Neighbour" effects. This is essentially the result of shared resources being consumed in extremis within a multi-tenant setup, meaning one VNF's resources are restricted by that of another VNF. One of the major shared resource bottlenecks is the central processor's Last Level Cache (LLC). [VEI2017] details a testbed which enables "Cache Allocation Technology" (CAT), so as to deterministically prioritize LLC resources between competing workloads. A number of CAT "Class of Service" (CoS) paradigms are explained for a range of service chain scenarios, involving virtual Firewall and virtual Router VNFs alongside a Noisy Neighbour VNF. Significant performance benefits are confirmed, bringing the performance of target VNFs in the presence of a LLC-hungry Noisy Neighbour, into alignment with the baseline scenario of a "Noise-Free" neighbouring VNF.

Even though NFV has been touted as the silver bullet for tackling a number of operator problems, the reality is that, in practice, it has proved hard to achieve the stable, predictable performance provided by hardware middleboxes, and so operators have essentially resorted to throwing money at the problem, deploying highly underutilized servers (e.g., one NF per CPU core) in order to guarantee high performance during peak periods and meet SLAs. In [YAS2017] we introduce HyperNF, a high performance NFV framework aimed at maximizing server performance when concurrently running large numbers of NFs. To achieve this, HyperNF implements hypercall-based virtual I/O, placing packet forwarding logic inside the hypervisor to significantly reduce I/O synchronization overheads. HyperNF *improves throughput by 10%-73% depending on the NF, is able to closely match resource allocation specifications (with deviations of only 3.5%), and to efficiently cope with changing traffic loads.*

Containers are in great demand because they are lightweight when compared to virtual machines. On the downside, containers offer weaker isolation than VMs, to the point where people run containers in virtual machines to achieve proper isolation. In [MAN2017]*, we examine whether there is indeed a strict tradeoff between isolation (VMs) and efficiency (containers). We find that VMs can be as nimble as containers, as long as they are small and the toolstack is fast enough. We achieve lightweight VMs by using unikernels for specialized applications and with Tinyx, a tool that enables creating tailor-made, trimmed-down Linux virtual machines. By themselves, lightweight virtual machines are not enough to ensure good performance since the virtualization control plane (the toolstack) becomes the performance bottleneck. We present LightVM, a new virtualization solution based on Xen that is optimized to offer fast boot-times regardless of the number of active VMs. LightVM features a complete redesign of Xen's control plane, transforming its centralized operation to a distributed one where interactions with the hypervisor are reduced to a minimum. LightVM can boot a VM in 2.3ms, comparable to fork/exec on Linux (1ms), and two orders of magnitude faster than Docker. LightVM can pack thousands of LightVM guests on modest hardware with memory and CPU usage comparable to that of processes.*

Unikernels are increasingly gaining traction in real-world deployments, especially for NFV and microservices, where their low footprint and high performance are especially beneficial. However, they still suffer from a lack of tools to support developers. In [SCH2017] *we present uniprof, a stack profiler that supports Xen unikernels on x86 and ARM and does not requires any code changes or instrumentation. Its high speed and low overhead (0.1% at 100 samples/s) makes it usable even in production environments, allowing the collection of realistic and highly credible data.*

To demonstrate that intelligent network services can be realized at high throughputs, in [GON2017a] *we present Net2Vec, a flexible high-performance platform that allows the execution of deep learning algorithms in the communication network. Net2Vec is able to capture data from the network at more than 60Gbps, transform it into meaningful tuples and apply predictions over the tuples in real time. This platform can be used for different purposes ranging from traffic classification to network performance analysis. Finally, we showcase the use of Net2Vec by implementing and testing a solution able to profile network users at line rate using traces coming from a real network. We show that the use of deep learning for this case outperforms the baseline method both in terms of accuracy and performance.* [GON2017b] *explores the use of Net2Vec for network data monetization use cases.*

2.3.1.2 Software Radio and DSP Virtualization

In contrast to previous work, SUPERFLUIDITY addresses the question of how to unify and abstract both the protocol stack as well as the hardware platform in order to allow (dynamic) instantiation of multiple vendor protocols. SUPERFLUIDITY attempts to open the radio hardware platform and enable third parties to design protocol stacks by defining a set of primitives and associated APIs enabling protocol modelling and construction according to dataflow paradigm. In addition to this, SUPERFLUIDITY will design the portable runtime dataflow engine enabling simultaneous seamless execution of multiple protocols. In order to cope with computational and real-time requirements of protocol stack processing, HW acceleration and its abstraction will complement the work in this area.

SDN and NFV are the embraced technologies for the backhauling of future 5G networks. VM and Docker container based deployments have received much attention. [ARA2017] presents the virtualization of a prototyped software defined RAN architecture by using VMs and Docker containers. In addition, it provides an analytical model for the generalized software defined RAN architecture with the practice of VM based and Docker container based implementations. Using measurements obtained from the two testbeds and the introduced queuing model, we compare their performances and analyse the two different architectures. Results verify the superiority of the Docker technology. Some observations from the behaviour of the testbeds are concluded for a better understanding of the VM and Docker container based technologies for the future development of 5G SDN controller.

2.3.1.3 High Performance, Software-based Packet Processing

SUPERFLUIDITY proposes to advance the state of this art by considering the use of highly heterogeneous hardware for network processing, to open up levels of performance and power consumption tuning that could only be seen in ASICs, while retaining the flexibility only afforded by software, a proposition particularly attractive for ISP clouds. This includes leveraging hardware just now hitting the market such as lowpower, small-sized microservers with architectures beyond the traditional x86 (e.g., MIPS64, ARM, ARM’s BIT.little, to name but a few) in order to severely reduce energy costs, especially for edge deployments.

The ability to automatically extract high performance out of commodity hardware remains largely an open research problem, and one that is only likely to receive greater attention as an increasing amount of network functionality moves to software.

SUPERFLUIDITY will aim at meeting the stringent requirements imposed by future 5G networks by designing and implementing a superfluid, converged network architecture that is location, hardware and time-independent. The work will push the boundaries of what is currently possible with virtualised, software-based packet processing (10-40Gb/s and higher, extremely fast service instantiation and migration in milliseconds, massive numbers of concurrent virtualised services on a single platform, significant power reductions, etc.). The goal is to bring the advantages of cloud and software-based networking to 5G networks so that services can be deployed whenever and wherever they are needed, and to leverage the availability of inexpensive, off-the-shelf hardware in the process. *In addition, SUPERFLUIDITY will contribute to the enhancement of underlying application delivery methods through novel mechanisms, thus facilitating seamless end-to-end user experience [MEK2016a, MEK2016b].*

SUPERFLUIDITY partners carefully addressed all issues related to core networking protocols, provided the desired backward compatibility required by future 5G deployments. TCP acceleration was a fundamental necessity for the whole consortium therefore a lightweight TCP proxy was designed and implemented, as described in [SIR2016]. This TCP proxy, called Miniproxy, was built on top of a specialized minimalistic cloud operating system, shows connection handling performance comparable to that of full-fledged GNU/Linux TCP proxy implementations, however its minimalistic footprint enables new use cases. Specifically, Miniproxy requires as little as 6 MB to run and boots in tens of milliseconds, enabling massive consolidation, on-the-fly instantiation and edge cloud computing scenarios. The benefits of Miniproxy were demonstrated by implementing and evaluating a TCP acceleration use case, clearly proving that this solution is capable of accelerating TCP's connection establishment and slow start phases. Miniproxy's very short boot times facilitate on-the-fly proxy boot, as the SYN packet of a TCP connection is first received.

In addition, Multipath TCP (MPTCP), the drop-in replacement for TCP that can use multiple paths in a single transport connection to improve reliability and throughput, was also evaluated. More specifically SUPERFLUIDITY partners analysed its fundamental ingredient, the congestion control algorithm that aims to efficiently allocate network resources across different multipath flows. The aforementioned analysis revealed that the MPTCP congestion controller may be gamed with a technique called policy drop, which adds artificial loss to all TCP and MPTCP traffic based on the client's access link speed and the connection RTT. This method redirects MPTCP traffic to alternative paths and may prove to be a simple way for network operators to regulate traffic [POP2016].

The starting point of [VAS2017] is the observation that all resources, including processors, main memory and storage are used, most often, to produce, consume or to transfer data, and can be viewed as special types of network "links" with different speeds. We want to explore what would happen if we do away with cloud-style reservations for in-network processing: Is it possible to adopt the best-effort per-packet service in the context of sharing other resources beyond bandwidth? In other words, we would like to rely on endpoint congestion control to share network processing resources in the same way it shares network capacity. We rely on experimentation to see whether traditional O/S *process scheduling is sufficient to share network processing. Our results show that sharing CPUs via traditional O/S mechanisms is suboptimal, as it leads to inefficient and unfair resource allocations in many situations. Next, we propose changes to the buffering discipline (endpoint congestion control) used by the network processes that remedy the problems identified in our experiments. We evaluate these changes in simulation, showing that they achieve both good fairness and resource allocation.*

Returning to the subject of TCP proxies, in [SIR2017] we focus on cases where the TCP proxy is only required during the initial phases of a network connection, becoming just a relay during the later stages, until the connection is finally closed. We try to answer the following question: Can established connections be offloaded from the TCP proxy? Our goal is to save precious resources by transparently removing the TCP proxy from the data path, when the proxy's operations are limited to relaying packets. Performing a TCP connection offload implies fulfilling two requirements. First, the offload should be performed without modifying the end-hosts of a connection, i.e., the client and the server. Second, the offload should not negatively affect the connection performance. Indeed, the main issue is in the need to join together two TCP connections with their specific per-connection states. To this end, our work provides a twofold contribution. First, we identify the minimum set of operations required to implement the joining of two TCP connections fulfilling the above requirements. Then, we show that such operations can be easily supported by already deployed infrastructure elements, such as P4 programmable switches. We implemented a proof-of-concept of our technique, named TCP proxy bypass, using PISCES, a P4 software switch based on OVS, and extending Miniproxy [SIR2016].

2.3.1.4 Automated KPI Mapping to Platform Metrics

For optimised service deployments (initial placement / rebalancing), a key question relates to what do we modify in our system to meet required KPIs? Where do we act and on which metrics? SUPERFLUIDITY will identify metrics, which mostly significantly influence service level KPIs (latency, throughput, etc.) with corresponding platform features (modification/action options). The approach taken will be based on an automated methodology which maps user defined service level KPIs to platform features and metrics, which have the most significant influence on the service performance. Correlation of service level KPI's to platform features/resources using a feature ranking approach built on classification/regression, statistical and ML techniques. Automatic KPI mapping will allow:

- *Better initial placement – enforce policies to ensure SLA compliance during placement*
- *Re-balancing – System observation based on most influential metrics (dimensionally reduce large metrics sets); trigger appropriate actuations (kill, migrate & tune) based on "influence knowledge insights".*
- *Capacity planning – Demonstration of the benefits of specification components/technologies within the infrastructure landscape which improve service performance.*

The applicability of KPI mapping for virtual infrastructure scaling in the use case of deploying a realistic video streaming service was demonstrated in [MEK2016a].

Fundamental to the Automated KPI Mapping approach is the ability to generate traffic workloads. In [PRE2017a] we present Hammer, a real-world, end-to-end network traffic simulator, capable of simulating complex and dynamic network, user and server behaviours. The focus of this tool is to primarily facilitate investigations related to product stability, for instance different aspects of capacity, longevity, memory leaks, cores and also handle customer content testing that will reveal the behaviour of the device under test in realistic network conditions. Hammer has a modular design, which offers excellent scalability, rendering the platform capable of being installed on commodity hardware. In addition, Hammer has a resource-savvy nature, thus requires limited computational resources to generate significant traffic loads. Instead of operating on a packet level, Hammer offers application-layer workload interaction along with inherent data plane acceleration, delivering brisk performance with unparalleled flexibility and ease of use. Our tests show that Hammers operation is linearly linked to the underlying hardware resources, however, even when the simulator is installed in a resource-bound environment, it can still deliver traffic loads that correspond to thousands of interconnected users each with real-world behaviour per session. To verify this in realistic scenarios, in [PRE2017b] we have partially benchmarked Hammer and demonstrated some of its capabilities on testing actual platforms deployed in production-like networking environments. Hammer was integrated to a complex testbed and proved capable of identifying the operational breaking point or the target node, as well as obtaining several metrics for its behaviour.

[MEK2018] demonstrates how the TALE methodology (see Section 2.1) can automatically profile virtualized media processing functions. Most media streaming services are composed by different virtualized processing functions such as encoding, packaging, encryption, content stitching etc. Deployment of these functions in the cloud is attractive as it enables flexibility in deployment options and resource allocation for the different functions. Yet, most of the time overprovisioning of cloud resources is necessary in order to meet demand variability. This can be costly, especially for large scale deployments. Prior art proposes resource allocation based on analytical models that minimize the costs of cloud deployments under a quality of service (QoS) constraint. However, these models do not sufficiently capture the underlying complexity of services composed of multiple processing functions. Instead, we introduce a novel methodology based on full-stack telemetry and machine learning to profile virtualized or cloud native media processing functions individually. The basis of the approach consists of investigating 4 categories of performance metrics: throughput, anomaly, latency and entropy (TALE) in offline (stress tests) and online setups using cloud telemetry. Machine learning is then used to profile the media processing function in the targeted cloud/NFV environment and to extract the most relevant cloud level Key Performance Indicators (KPIs) that relate to the final perceived quality and known client side performance indicators. The results enable more efficient

monitoring, as only KPI related metrics need to be collected, stored and analyzed, reducing the storage and communication footprints by over 85 %. In addition a detailed overview of the functions behavior was obtained, enabling optimized initial configuration and deployment, and more finegrained dynamic online resource allocation reducing overprovisioning and avoiding function collapse. We further highlight the next steps towards cloud native carrier grade virtualized processing functions relevant for future network architectures such as in emerging 5G architectures.

2.3.2 Network Services Decomposition and Programmability

SUPERFLUIDITY will devise programming abstractions specifically targeted to 5G functions. The API design work will address three programming levels: service, function, and processing levels, and will attempt to maximise viability by reusing existing standard (when applicable) or research community best practices. Work will on one side target the definition of 5G specific actions and events, and on the other side will address the specification of the constructs needed to combine and orchestrate a desired execution of such actions (conditioned on the arrival of events). Particularly promising and forward-looking is SUPERFLUIDITY's approach of combining block-based composition abstractions (such as those exploited in Click routers, in some software defined radio architectures, or emerging in the ETSI NFV work on service chaining) with event-driven programming paradigms such as basic match/action based approaches or more powerful stateful abstractions based on extended finite state machines *[ETT2016]*.

With the introduction of SDN in large-scale IP provider networks being still an open issue with many different solutions constantly being suggested, SUPERFLUIDITY researchers proposed a hybrid approach that allows the coexistence of traditional IP routing with SDN based forwarding within the same provider domain. The solution is called OSHI – Open Source Hybrid IP/SDN networking and is fully implemented by combining and extending Open Source software. The OSHI system architecture is described in [NSM2016] and includes the design and implementation of advanced services like Pseudo Wires and Virtual Switches. In addition, a set of Open Source management tools for the emulation of the proposed solution using either the Mininet emulator or distributed physical testbeds was demonstrated, which include an extensible web-based graphical topology designer, providing different layered network "views" (e.g. from physical links to service relationships among nodes). The suite is capable of validating an input topology, deploy it in an automatic manner over a Mininet emulator or a distributed SDN testbed and allow access to emulated nodes by opening consoles in the web GUI, along with the necessary tools for evaluating the performance of all deployed nodes.

In [SCH2016] we define the notion of composition for SDN applications and show the theoretical and practical approaches to composition in SDNs, explaining the challenges associated with it. We explore *the feasibility of OpenFlow as an Application Programming Interface (API) for a composition engine and argue that its design as Southbound controller interface makes it unsuitable for this task.*

Provided that Network Service complexity will significantly increase in the foreseeable future, all network function outsourcing and efficient deployment techniques will render 5G infrastructure fast, secure and content independent. For instance, a scalable system that will allow cloud service providers to privately compute network functions on behalf of a client in such a way that the cloud does not become aware of the underlying policies, thus providing strong security guarantees in the honest-but-curious model based on cryptographic secret sharing, is definitely a welcome addition to the proposed architecture in terms of security, network resource utilisation and security. Such a solution, called SplitBox was presented in [ASG2016] by first describing the abstract functionality model the implementation was based onto, and then by executing a firewall test case showing the solutions capabilities of achieving throughput of over 2 Gbps with 1kB-sized packets on average, traversing up to 60 firewall rules, numbers way beyond competition or existing implementations.

Programmable dataplanes are emerging as a disruptive technology to implement network function virtualization in an SDN environment. This technology can be further enhanced by using data plane abstraction with stateful processing. In [BON2017a] we focus on real world use cases with stateful forwarding requirements to validate Open Packet Processor (OPP) data plane abstraction. We first demonstrate the suitability of OPP for implementing complex stateful network functions by providing the detailed implementation of three use cases, which were demo'ed in [BON2017b]. Second, we assess the scalability of the use case implementations in the context of a datacenter deployment.

[DAR2017] is an extensive survey we published on IEEE Communications Surveys and Tutorials that discusses the security implications of stateful SDN planes. SDN emerged as an attempt to introduce network innovations faster, and to radically simplify and automate the management of large networks. SDN traditionally leverages OpenFlow as device-level abstraction. Since OpenFlow permits the programmer to "just" abstract a static flow-table, any stateful control and processing intelligence is necessarily delegated to the network controller. Motivated by the latency and signaling overhead that comes along with such a two-tiered SDN programming model, several works have proposed innovative switch-level (data plane) programming abstractions capable to deploy some smartness directly inside the network switches, e.g., in the form of localized stateful flow processing. Furthermore, the possible inclusion of states and state maintenance primitives inside the switches is currently being debated in the OpenFlow standardization community itself. In [DAR2017], after having provided the reader with a background on such emerging stateful SDN data plane proposals, we focus our attention on the security implications that data plane programmability brings about. Also via the identification of potential attack scenarios, we specifically highlight possible vulnerabilities specific to

stateful in-switch processing (including denial of service and saturation attacks), which we believe should be carefully taken into consideration in the ongoing design of current and future proposals for stateful SDN data planes.

To highlight the benefits of the SUPERFLUIDITY architecture model, [CHI2017c] considers the problem of evaluating the performance of a 5G network based on RFBs, as proposed by the project. RFBs allow a high level of flexibility, agility, portability and high performance. After formally modelling the RFB entities and the network physical nodes, we optimally formulate the problem of maximizing different Key Performance Indicators (KPIs) on a RFB-based network architecture, in which the RFBs are shared among the nodes, and deployed only where and when they are really needed. Our results, obtained by solving the proposed optimization problem over a simple yet representative scenario, show that the network can be managed in a very efficient way.

In [SHO2017], we start from a 5G Superfluid Network, as an outcome of this project. The network apparently exploits the concept of Reusable Functional Block (RFB), a virtual resource that can be deployed on top of 5G physical nodes. Then we target the management of the RFBs in a Superfluid network to deliver a high definition video to the users. Specifically, we design an efficient algorithm, called P5G, which is based on Particle Swarm Optimization (PSO). Our P5G solution targets different Key Performance Indicators (KPIs), including the maximization of user throughput, or the minimization of the number of used 5G nodes. Results, obtained over a representative scenario, show that P5G is able to wisely manage the RFBs, while always guaranteeing a large throughput to the users.

Optimal design of 5G Superfluid Networks is also discussed in [CHI2018]. The core of the project is the definition of a 5G converged architecture based on virtual entities, Reusable Functional Blocks (RFBs), which can be run on different HW and SW execution environments. The exploitation of RFBs allows to achieve the required level of flexibility required by 5G. After optimally formulating the problem of minimizing the total installation costs of a SuperFluid Network composed of RFBs and physical 5G nodes, we propose a new algorithm, called SFDA, to practically tackle the problem. Our results, obtained over a representative case study, show that SFDA is able to solve the problem in a reasonable amount of time, returning solutions very close to the optimum. In addition, we clearly show the trade-offs that emerge between the need of providing a service level to users (in terms of downlink traffic or coverage) and the total costs incurred to install the elements of the network.

Another important 5G KPI is energy efficiency. [TAJ2018] discusses the subject of energy efficient path allocation for Service Function Chaining (SFC). SFC is a service deployment concept that promises cost efficiency and increases flexibility for computer networks. On the other hand, SDN provides a powerful infrastructure to implement SFC. In [TAJ2018], we mathematically formulate the SFC problem in SDN-

based networks. In this way, the energy consumption of the network is minimized while the traffic congestion is controlled through network reconfiguration. Additionally, a low complex heuristic algorithm is proposed to find a near-optimal solution for the mentioned problem. Simulation results show that the proposed heuristic reconfigures the network in a way that the energy consumption is near-optimal while the SFC requirements are met. Besides, the computational complexity is very low which makes it applicable for real-world networks.

In [SAL2017a] we present RDCL 3D, a "model agnostic" web framework for the design and composition of NFV services and components. The framework allows editing and validating the descriptors of services and components both textually and graphically and supports the interaction with external orchestrators or with deployment and execution environments. RDCL 3D is open source and designed with a modular approach, allowing developers to "plug in" the support for new models. We describe several advances with respect to the NFV state of the art, which have been implemented with RDCL 3D. We have integrated in the platform the latest ETSI NFV ISG model specifications for which no parsers/validators were available. We have also included in the platform the support for OASIS TOSCA models, reusing existing parsers. Then we have considered the modelling of components in a modular software router (Click), which goes beyond the traditional scope of NFV. We have further developed this approach by combining traditional NFV components (VNFs) and Click elements in a single model. Finally, we have considered the support of this solution using the Unikernels virtualization technology.

In the context of the emergent SDN paradigm, the attention is mostly directed to the evolution of control protocols and networking functionalities. However, network professionals also need the right tools to reach the same level—and beyond—of monitoring and control they have in traditional networks. Current SDN tools are developed on an ad-hoc basis, for specific SDN frameworks, while production environments demand standard platforms and easy integration. [ROJ2018] is an extensive survey, published on ACM Computing Surveys, that aims to foster the definition of the next generation SDN management framework by providing the readers a thorough overview of existing SDN tools and main research directions.

2.3.3 Cloud RAN and Mobile Edge Computing

Beyond the current vision of a static RAN function fully located in one "edge computing" place, SUPERFLUIDITY will support the ability to modularly "hot" replace eNB functions (such as scheduling) and to permit migration of such functions between edge clouds and the antenna subsystem, so as to balance algorithmic complexity with fronthaul capacity. SUPERFLUIDITY will also transcend current Mobile Edge Computing vision where non-RAN functions (local caching, CDN, ...) are envisaged to be co-located only at the eNB by enabling their migration between the RRH and the edge cloud, to maximise their performance. *In addition, SUPERFLUIDITY could exploit the Mobile Edge Computing* *architecture to design, implement, deploy and evaluate additional functional blocks for monitoring purposes [TSE2016].*

SUPERFLUIDITY also advances the state of the art by decomposing video streaming functions to exploit the combined core and edge architecture. For example, edge (late) transmuxing [MEK2016b] was developed, which showed improved video streaming performance by moving the trans multiplexing operation to the edge, while keeping the core content at the core. Further experiments planned with DRM and Personalization at the edge are expected to achieve similar improvements. SUPERFLUIDITY develops new communication strategies between edge and core to realize these improvements.

Recent advances in point cloud capture and applications in VR/AR sparked new interests in the point cloud data compression. Point Clouds are often organized and compressed with octree based structures. The octree subdivision sequence is often serialized in a sequence of bytes that are subsequently entropy encoded using range coding, arithmetic coding or other methods. Such octree based algorithms are efficient only up to a certain level of detail as they have an exponential run-time in the number of subdivision levels. In addition, the compression efficiency diminishes when the number of subdivision levels increases. In [AIN2016] we present an alternative enhancement layer to the coarse octree coded point cloud. In this case, the base layer of the point cloud is coded in known octree based fashion, but the higher level of details are coded in a different way in an enhancement layer bit-stream. The enhancement layer coding method takes the distribution of the points into account and projects points to geometric primitives, i.e. planes. It then stores residuals and applies entropy encoding with a learning based technique. The plane projection method is used for both geometry compression and color attribute compression. For color coding the method is used to enable efficient raster scanning of the color attributes on the plane to map them to an image grid. Results show that both improved compression performance and faster run-times are achieved for geometry and color attribute compression in point clouds.

In [MEK2016c] we present MP3DG-PCC, an open source framework for design, implementation and evaluation of point cloud compression algorithms. The framework includes objective quality metrics, lossy and lossless anchor codecs, and a test bench for consistent comparative evaluation. The framework and proposed methodology is in use for the development of an international point cloud compression standard in MPEG. In addition, the library is integrated with the popular point cloud library, making a large number of point cloud processing available and aligning the work with the broader open source community. [MEK2017] presents a methodology for performance assessment that was established in the MPEG point cloud compression activity in a collaboration of leading companies and academic institutes. It includes selected datasets, objective metrics and target coding rates corresponding to distinguishable subjective qualities and different use cases. The performance assessment methodology is applied to a compression framework on a large dataset developing an *anchor performance benchmark. The benchmark results and methodology are used in the recent MPEG Call for proposals on point cloud coding to consistently compare and evaluate technologies.*

[WAG2017] presents Unified Remix, a solution for adaptive bit-rate streaming of video presentations with inserted or edited content. The solution addresses three important challenges encountered when streaming personalized media presentations. First, it reduces vulnerability to ad blocking technologies and client-side playback deviations encountered when using manifest manipulation based methods. Second, it reduces storage and computational costs associated with alternative server side solutions such as brute force re-encoding or duplicate storage towards levels comparable to linear video streaming (VoD or Live). Third, it handles the multi-source, multi-DRM and multiprotocol aspects for modern video streaming natively in the workflow. The solution is based on a combination of existing proven streaming technologies such as Unified Origin and newly designed components such as the Remix MPEG-4 module. The framework uses standardized technologies such as MPEG-4 ISOBMFF, SMIL and MPEG-DASH. The components work together in a micro service architecture enabling flexible deployment using a (container) orchestration framework on premises or in the cloud. The solution is demonstrated in two use cases: content pre-/post/mid roll and Live Archive to VoD conversion. As many use cases can be implemented based upon Unified Remix, we envision it as a key component of professional video streaming platforms.

In [SAL2017b], we first introduce the architecture for dynamic deployment and composition of virtual functions proposed by the Superfluidity project. Then we consider a case study based on a typical 5G scenario. In particular, we detail the design and implementation of a Video Streaming service exploiting Mobile Edge Computing (MEC) functionalities. The analysis of the case study provide an assessment on what can be achieved with current technologies and gives a first confirmation of the validity of the proposed approach. Finally, we identify future directions of work towards the realization of a Superfluid softwarized network.

In [ESP2017], we present a model for the complete edge function onloading problem, which consists of three main phases: (1) Cyber foraging, which involves discovery of resources monitoring the state of edge resources, (2) edge function mapping, which involves matching requests to available resources, and (3) allocation, which involves assigning resources to mappings. Using optimization theory, we show how these three phases are tightly connected, and how the wide spectrum of existing solutions that either solve a particular phase, or jointly solve two of the phases (along with their interactions), are incomplete and may lead to inefficiencies. Moreover, with extensive simulation experiments we demonstrate that joint optimization of all three phases enables the edge network to host a larger set of constrained edge function requests.

One of the most challenging KPIs defined by the ITU-T (IMT-2020) for 5G is "latency below 1 ms". This is a key requirement for an emerging era of new applications, such as virtual and augmented reality, video analytics, or Industry 4.0. Those applications require low latency to enable the so-called tactile Internet. However, relying on large centralized datacenters as today, this requirement defies the laws of physics. The solution is edge computing, a technology aiming to deploy applications at the edge of the network, closer to end users. Altice Labs has been contributing to the standardization of this technology from the beginning, and is currently prototyping a fully functional and ETSI standards-compliant solution. The first version of this prototype has been already released and a proof-of-concept (PoC) has been deployed at the Altice PT Labs. [PAR2018] describes the edge computing technology, the prototype under development, the PoC integration and the results of the preliminary evaluation already performed.

2.3.4 Automated Security and Correctness

SUPERFLUIDITY will provide a two-pronged, complementary approach to security. First, it will go beyond the state of the art, providing a pre-deployment checking system that will ensure that virtualised network services do not negatively affect the network nor other tenants; unlike approaches in the literature, the system will be both scalable and stateful, able to model most types of services. Second, SUPERFLUIDITY will implement a post-deployment system that will learn the behaviour of traffic and detect any anomalies, thus providing a further security mechanism in cases where the checking system does not have information about the processing performed by a network function, or when static analysis is inaccurate.

The principle mission of Group Testing (GT) is to identify a small subset of "defective" items from a large population, by grouping items into minimum number of test pools. The test outcome of a pool is positive if it contains at least one defective item, and is negative otherwise. GT algorithms are utilized in many applications, and the privacy regarding the status of the items, namely, defective or not, is a critical issue. In this paper, we consider a scenario where there is an eavesdropper (Eve) which is able to observe a subset of the GT outcomes (pools). In [COH2016] we propose a new non-adaptive Secure Group Testing (SGT) algorithm based on information theoretic principles, which keeps the eavesdropper ignorant regarding the items status.

In [KAM2016], we study the secrecy rate and outage probability in Multiple-Input-Single-Output (MISO) Gaussian wiretap channels at the limit of a large number of legitimate users and eavesdroppers. In particular, we analyse the asymptotic achievable secrecy rates and outage, when only statistical knowledge on the wiretap channels is available to the transmitter. The analysis provides exact expressions for the reduction in the secrecy rate as the number of eavesdroppers grows, compared to the boost in the secrecy rate as the number of legitimate users grows.

2.3.4.1 Correctness Checking using Symbolic Execution

Next generation network deployments will have an elevated level of complexity due to the increased functionality they are designed to support along with their user-oriented nature. To tackle potential issues, new tools need to be introduced, some of which are already under investigation by the SUPERFLUIDITY partners. Solutions based on Symbolic Execution are also implemented by the consortium members, for instance the SymNet, a network static analysis tool described in [STO2016a]*. SymNet quickly analyses networks by injecting symbolic packets and tracing their path through the network. The key novelty behind SymNet is SEFL, a language specifically designed for network processing which can be categorised as symbolic-execution friendly. SymNet is easy to use: parsers that automatically generate SEFL models from router and switch tables, firewall configurations and arbitrary Click modular router configurations were developed. Most of the underlying models are exact and have optimal branching factor. A testing tool that checks SEFL models resemblance to the real implementation was also integrated rendering SymNet able of checking networks containing routers with hundreds of thousands of prefixes and NATs in seconds, while ensuring packet header memory-safety and capturing network functionality such as dynamic tunnelling, stateful processing and encryption. SymNet was used to debug middlebox interactions documented in the literature, to check numerous existing networks and troubleshoot the basic networking component of OpenStack, Neutron* [STO2016b]*. Results show that symbolic execution is fast and more accurate than existing static analysis tools. Especially in the OpenStack Neutron case, despite being work in progress, revealed a propagation delay bug and allowed tenants to deploy networks with relative ease.*

2.4 Innovation Potential

SUPERFLUIDITY's ambition is to aid mobile network operators to meet the demanding requirements of 5G while at the same time develop favourable market conditions, which will create sustainable business for all players in the value chain. In an environment of slowing market expansion rates and declining revenues, where end users request personalised services, better performance, better user experience, businesses need to get more information about their consumers, easier and secured access to devices and greater flexibility for provisioning new services. There is a key role to play for Equipment providers, Service Providers and IT players together to make this a reality by providing converged IT and Network infrastructure.

SUPERFLUIDITY can provide this, thanks to the unique blend of Telco and IT players in its consortium.

In this context, SUPERFLUIDITY helps to create a new ecosystem and the opportunity for all players to collaborate and develop new business models they can each benefit from. More in detail, SUPERFLUIDITY innovation is to:

- Enable Mobile Network Operators (MNOs) to rapidly deploy new services for consumer and enterprise business segments which can help them differentiate their service portfolio and to adapt them instantaneously (time-independence) to varying traffic demands by localising processing (location-independence), providing users with the best quality of experience constantly;
- Boost the capabilities of software and application providers by allowing them to develop applications that can take advantage of the information (through specific APIs) on radio network capabilities and conditions available at the base station without the need to worry about hardware (hardware-independence), performance and scaling issue (scale-independence).

Regarding specific market segments, the main innovation arising from a careful analysis of SUPERFLUIDITY innovation potential is that nobody in the market so far can claim to have even roadmapped such a vision (convergence of IT and Network infrastructure all the way from the core towards the access, open APIs for exploiting radio network information, near-instantaneous relocation of services, independence from hardware components, etc.).

Below we report an extract of such an analysis to substantiate our statements:

- "LTE Network Infrastructure"
    - Leaders (per Gartner, April 2015): Ericsson, Nokia, Alcatel-Lucent, Huawei (NEC is in the Visionaries)
    - SUPERFLUIDITY innovation: the converged architecture concept of SUPERFLUIDITY will guide network equipment vendors to evolve their infrastructure products to 5G.

- "Small Cell Equipment"
    - Leaders (per Gartner, October 2015): Huawei, Alcatel-Lucent, Nokia, Ericsson, Cisco, ZTE (NEC is in the Challengers)
    - SUPERFLUIDITY innovation: small cells will play a very important role in fulfilling 5G requirements. The advances in CRAN and MEC, as also contributed by SUPERFLUIDITY, will evolve these products through the adoption of softwarisation and cloudification.

- "x86 Server Virtualization Infrastructure"
    - Leaders (per Gartner, July 2015): VMware, Microsoft (Red Hat and Citrix are in the Niche Players)
    - SUPERFLUIDITY innovation: even though there are no plans to improve hypervisors, the adoption of minified VMs, unikernels, containers & specialised operating systems, and the reduction of orchestration system overhead, will bring those to the next level.

- "IT Services for Communications Service Providers, Worldwide"
    - Leaders (per Gartner, July 2015): Accenture, IBM
    - SUPERFLUIDITY innovation: provide more powerful platform that has the ability to homogeneously manage the network and its functions as a cloud platform across heterogeneous infrastructures all the way from the core to the access and vice-versa.
- "Cloud Infrastructure as a Service, Worldwide"
    - Leaders (per Gartner, May 2015): Amazon Web Services, Microsoft / Azure (Google, CenturyLink, VMware, IBM / SoftLayer are in the Visionaries)
    - SUPERFLUIDITY innovation: enable more powerful private and hybrid cloud offerings by empowering solutions to offer, transparently and with performance guarantees, the best of both bare-metal managed hosting and virtualised cloud environments spanning all the way towards the access of the network.

- "Public Cloud Infrastructure as a Service, Worldwide"
    - Critical Capabilities (per Gartner, October 2015)
        - Application Development: Amazon Web Services, Microsoft, Virtustream
        - Batch Computing: Amazon Web Services, Microsoft, Google, Rackspace
        - Cloud-Native Applications: Amazon Web Services, Microsoft, Google
        - General Business Applications: Amazon Web Services, Virtustream, Microsoft
    - SUPERFLUIDITY innovation: same as above

- "Public Cloud Storage Services"
    - Leaders (per Gartner, June 2015): Amazon Web Services, Microsoft, Google
    - SUPERFLUIDITY innovation: same as above

- "Cloud-Enabled Managed Hosting, Europe"
    - Leaders (per Gartner, June 2015): Rackspace, Interoute, BT Global Services, Claranet, Colt, CenturyLink, Verizon (Telefonica is in the Niche Players)
    - SUPERFLUIDITY innovation: same as above; OnApp is a large player in this segment.
- "Managed Hybrid Cloud Hosting, Europe"
    - Leaders (per Gartner, June 2016): Interoute, BT, Rackspace, Claranet, CenturyLink
    - SUPERFLUIDITY innovation: same as above
- "Data Center Networking"
    - Leaders (per Gartner, May 2016): Cisco, Arista (Dell and VMware are in the Visionaries, Brocade, Juniper and HPE in the Challengers, while NEC is in the Niche Players)
    - SUPERFLUIDITY innovation: will significantly advance the NFV and SDN state of the art, through introducing recursive network service (NS) modelling and decomposition, and more powerful/stateful SDN abstractions, based on extended finite state machines.
- "Network Equipment and Solutions Providers"
    - Leaders (per Gartner, August 2013): Alcatel-Lucent, Cisco, Ericsson, NEC, Nokia.
    - SUPERFLUIDITY innovation: offering a solution that spans over heterogeneous hardware architectures, offering performance- and scale-independent solutions in the cloud for networking deployments that are closely following the "on-demand" cloud computing innovations.

3 Exploitation Activities and Plans

Activities in this area aim to fulfil the Market Vision of the SUPERFLUIDITY project (please refer to Section 3.3 of deliverable *D1.3 Project Vision and Roadmap, v2*).

Section 3.1 outlines the exploitation plans of the individual partners. These were initialised from the Description of Action (DoA), but we revised them per the strategic or vision shifts of each partner.

Section 3.2 summarises the exploitation plan of the consortium as a whole. Other than to reflect consolidations (M&A), the overall plan has remained largely unchanged, which is a testimony to the insightfulness of project objectives, at the time of project definition.

As definition and implementation of 5G accelerates and the associated market segments continue to "pivot", we will keep these exploitation plans aligned. Such changes were reported, both internally, as reports of the Innovation Coordinator to the Management Board, as well as to the Project Officers, in the context of the regular reports and annual reviews.

3.1 Individual Partners

The *updated* exploitation plans of individual SUPERFLUIDITY partners are reported in Table 3 below. *Due to the fact EBLINK is no longer a partner of the project, the respective exploitation activities have been undertaken by NOKIA FR, through a new Linked Third Party named "Nokia xHaul Business Line".*

PARTNER	EXPLOITATION PLANS
NOKIA FR	The SUPERFLUIDITY project brings significant contributions to the different Nokia solutions and products, from the application cloud management to the wireless part. The work conducted in SUPERFLUIDITY on Cloud RAN is highly valuable since it allows Nokia Bell-Labs to analyse and identify the right technological choice for the virtualization, hardware acceleration, service chaining, orchestration tool, in an advanced stage preparing the next iteration of the Cloud RAN product. *In addition to, the work conducted on the front haul is of high interest to our Business Line which is now working as Linked Third Party and benefits from the platform feedback developed in WP7.*

	SUPERFLUIDITY helps Nokia Bell-Labs to realise and prototype the vision of a cloud-based virtualised wireless access. This exploitation will facilitate the shift of additional applications from the cloud to the remote antenna and vice-versa, increasing the resource optimization at the cloud infrastructure. Most importantly, such a solution will allow the virtualization of RAN components and thus the convergence of all access (fixed and wireless) services within a single framework. *All components (RFBs) will be deployed in the EDGE Cloud similar to what we are implementing in WP7 related to side-by-side VM/container deployment.* In addition, the collaboration with partners of the SUPERFLUIDITY project, allows us to enhance our cloud RAN via RFB principle (component based), integrate the Mobile Edge Computing and evaluate the benefits that we can extract for a Cloud RAN performance. This step is very important to prepare validate the Cloud RAN design and facilitate the discussion with NOKIA Mobile Network Business Unit. *This is re-enforced by joining Nokia xHaul Business Line in the project.* Furthermore, the project results may improve multi-tenancy and multi-operator capability of future 5G Nokia products (e.g. Netact, Airframe, Next Cloud packet Core...) through the use of cloud, SDN and NFV technologies. It is worth to highlight that the next Core product from Nokia will be based on micro-services decomposition (see the recent press release https://networks.nokia.com/products/mobility-manager). The different area of investigation in SUPERFLUIDITY (related to Infra-structure as a service, Platform as a service, Virtualization, Orchestration QoE aware) give us good inputs to be considered in Nokia product, as cloud management product, MEC product, … *Furthermore, Nokia uses results from Superfluidity in direct customer presentations and discussions with customers to shape the 5G mobile network solution. Nokia Bell Labs has a strong contact to Nokia business lines and provides input to the definition and setup of the 5G product roadmap. In particular, Nokia Bell Labs contributed specific input on Cloudification of RAN, the selection of virtualization technology, the support of distributed clouds to the Nokia Mobile Networks business group. Also, the microwave Wireless fronthaul technology, the Nokia xHaul product, benefit from all the test and analysis conducted in WP7 through the use of different RAN split which allows to improve the product and as a consequence gain more markets and business. Nokia will further explore the application of RFB technology in Next Generation Platform as a Service project, led by Nokia Bell-Labs. Furthermore, Superfluidity results will be applied to optimize the roadmap of moving some products using VM technology toward container and service based architecture.*
NOKIA IL	NOKIA IL is targeting to exploit SUPERFLUIDITY's results to evolve NOKIA's portfolio with respect to NFV and MEC.

	First, extending NOKIA's NFV orchestration to better support C-RAN and MEC application is an immediate outcome of SUPERFLUIDITY. Then, SLA based placement across different POPs from the cloud edge to the core network is yet another enhancement to Cloudband VNF manager. It is expected, as envisioned by SUPERFLUIDITY, that future network functions would be composed of microservices that could be rapidly instantiated, scaled or migrated. Providing the platform for those microservices, e.g., via containers and uni-kernels, is in the plans of Cloudband Infrastructure software. Additionally, extending the orchestration and management to support life cycle management in such a dynamic environment is in the roadmap of Cloudband Network director as well as Cloudband Application Manager.
BT	BT believes that in the next 2-3 years the worlds of cloud computing and virtualised network functions are coming together into a completely new model for telecommunications combining cloud services, network function virtualisation, and software defined networks. Entire enterprises will be able to run on virtualised servers, virtualised network functions, and virtualised network control services. The BT Vision is described by Howard Watson, our CEO of BT Technology, Service & Operations, at https://player.vimeo.com/video/166341115 (starts at 14 minutes) As part of this convergence, NFV has a critical role. We have already launched our first commercial service that uses NFV (Cloud Connect Intelligence), with more in the pipeline. The fast and flexible deployment and modification of services is what is most attractive about NFV. We are feeding the results from SUPERFLUIDITY into our internal OSS/BSS developments. Several technologies researched in SUPERFLUIDITY are of specific interest to BT. From an implementation point of view these include the work on extended finite state machines (since these hold the promise of having full defined and understood behaviour) and lightweight virtualisation techniques, such as containers. The approach/tool to verify composition also addresses an important issue. Once these techniques are more mature, we will consider whether they are suitable for our network. BT is also interested in the various demonstrations. For instance, currently we have an issue with CPE in the branches of, say, a bank – because site visits are expensive and hard to organise. Virtualisation offers the prospect of simplifying the equipment in the branch, by moving the more complex and maintenance-prone functions into a central site. *BT is currently building its own internal 5G testbed, covering fixed and mobile networks. BT intends it to use the same equipment as in, or planned for, the real network, yet carry out experiments with new ideas (whether from SUPERFLUIDITY, other H2020 projects and internal research) such as: micro-services (RFBs), NFV virtualisation, containers and lightweight VMs,*

<table>
<tr><td></td><td>slicing, automation for network management, hierarchical service provision, and so on. In addition, some of the Superfluidity ideas (such as RFBs) are being further taken on in the NGPaaS project.</td></tr>
<tr><td>CITRIX</td><td>CITRIX is targeting to exploit SUPERFLUIDITY's results to evolve its products and corresponding solutions for Communications Service Providers, which have been migrated from the ByteMobile Adaptive Traffic Manager (T-3100 series) to the NetScaler Mobile Gateway (T-1000 series) product line (see https://www.citrix.com/solutions/communications-service-providers.html).

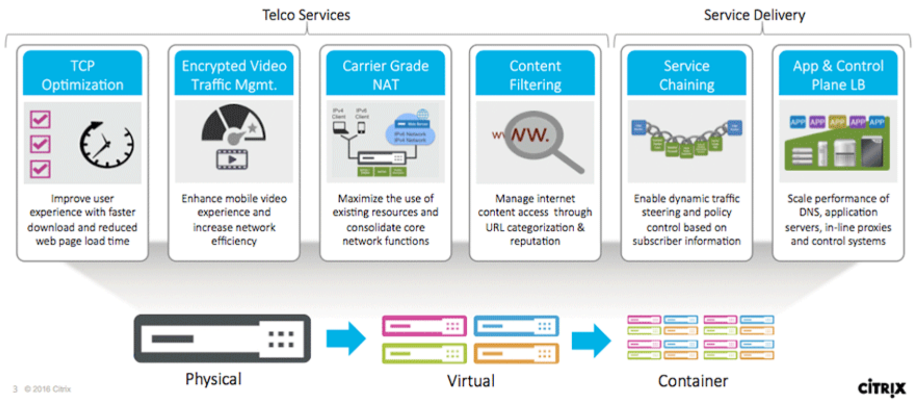

SUPERFLUIDITY's results will allow CITRIX to:

<u>Extend the applicability of NetScaler product family from the core network towards the wireless access by co-locating such (possibly virtualised) products next to wireless nodes (e.g., base stations) and taking advantage of radio network information to better optimise the traffic</u>. This will leverage the MEC infrastructure that is implemented as part of SUPERFLUIDITY. Feasibility of this will be proven through a PoC that CITRIX has planned to execute in the context of the project.
<u>Improve the performance of the virtualised appliances</u> (e.g., load balancing and application delivery) in order to offer a scalable service chaining orchestration of elementary functionalities with guaranteed SLA enforcement. This will leverage SUPERFLUIDITY innovations in the area of accelerating data-plane applications, particularly ones that require scalable load balancing and service function chaining (SFC).
</td></tr>
<tr><td>INTEL</td><td>NFV and SDN are foundational technologies in the realisation of 5G networks that are elastic, programmable and dynamically manageable. Areas such as enhanced platform awareness (EPA), workload affinity for platform features, intelligence orchestration and approaches to resource description models will play a critical role in the deployment of NFV and SDN for 5G. Research in WPs 4-6 into these topics within SUPERFLUIDITY mapped directly to Intel's internal research agenda. To ensure convergence between technologies being developed by Intel Business Units (BUs) and external customer needs, we aligned the efforts of our activities within SUPERFLUIDITY with our internal research framework</td></tr>
</table>

	Apex Lake which is focused on realising intelligent orchestration proof of concepts. From a commercial perspective, Intel's exploitation plans in SUPERFLUIDITY lie in the realisation that convergence of fixed and mobile technologies presents the opportunity for sharing of infrastructure resources for efficient service delivery. Intel is bring SUPERFLUIDITY results to the community to enable an efficient and highly manageable infrastructure landscape, making use of Standard High Volume (SHV) elements to empower XaaS (Platform or Infrastructure as a Service) solutions with resource pooling, on-demand scaling of capacity and allocation of resources proportionate to revenue. We completed an internal transfer of the automated methodologies within the project such as automated KPI mapping and automated optimisation of deployment templates based on the combination of telemetry and analytical approaches to contribute significantly in the realisation of a scalable, intelligent and automated orchestration platform. This work is will be exploited by our internal BUs to solve customer problems and challenges. *The results from the project were used to inform the formulation of a new a research agenda within Intel Labs Europe that started in 2018 which is focused towards the network edge.* Intel is actively engaged in both open standards and commercial software enablement. Intel has a strong involvement in a number of key open source projects aligned to NFV and SDN including OpenStack, OpenDaylight, Open vSwitch, OpenFlow and the Data Plane Development Kit (DPDK). In addition, Intel is a leading member of the Mobile Edge Computing (MEC) industry initiative, which is focused on opening up the radio access interface for services and applications in the radio access network (RAN). Intel's team located in Shannon, Ireland are the lead engineering group for much of the activity in SDN/NFV, and Intel will use SUPERFLUIDITY as proof point for open source technologies, MEC and the value of an open source ecosystem to ensure interoperability for 5G deployments and associated use cases.
NEC	NEC is targeting to exploit SUPERFLUIDITY's results to evolve its solution portfolio in three main areas: • Small cell solution (both outdoor and in-building) (http://www.nec.com/en/global/solutions/nsp/sc/index.html) • LTE-Advanced high-capacity base station equipment (http://www.nec.com/en/press/201302/global_20130227_04.html) • Cloud solution (http://www.nec.com/en/global/solutions/cloud/index.html?)

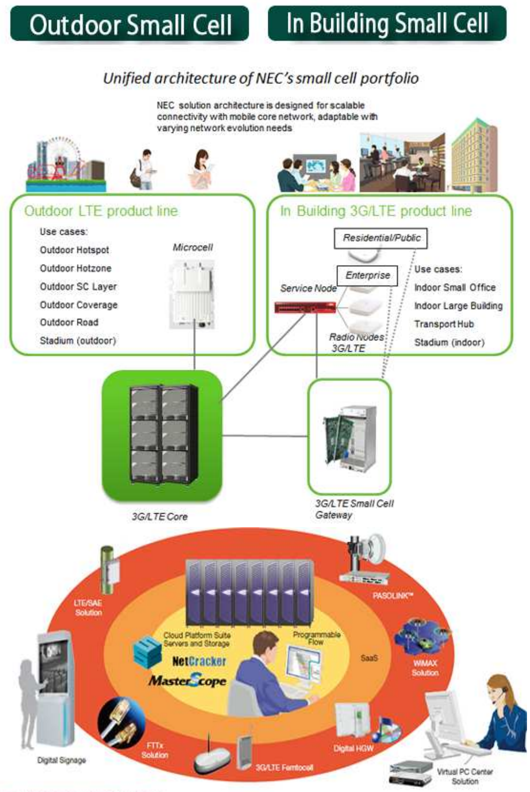

Figure 5. NEC's Solutions for Carrier Cloud Services

In outdoor scenarios, NEC will target the technology transfer to its business units responsible for the "New Last Mile" concept: the one-stop-shop to help mobile network operators deploy the next layer of LTE capacity and coverage in a rapid, economical and integrated way and keep pace with skyrocketing data demand. Such solution is designed to simplify the end-to-end process of delivering the LTE brand promise in complex 3D urban landscapes, VIP customer sites outside the current macro footprint and rural areas in line with their spectrum licensing obligations. The initiative leverages NEC's leadership in open, multivendor HetNet technology and the most comprehensive small cell and backhaul portfolio in the market.

In in-building scenarios, we target a technology transfer to NEC's Enterprise RAN (E-RAN) solution that already utilise innovative multimode LTE and 3G coverage for enterprise internal phone network.

As for LTE-Advanced High-Capacity Base Station Equipment, NEC is planning to evolve its base station from a system where specialist processing is done to a modular, open solution integrated in an ecosystem of changeable components as the "evolution of the base station concept".

Overall, NEC's exploitation plan for SUPERFLUIDITY is to integrate the three above areas in a horizontal manner and provide a true end-to-end solution with mobile edge to cloud cooperation (seamless and quality-aware migration of service and applications from radio head to cloud and vice-versa).

ONAPP	OnApp intend to exploit SUPERFLUIDITY's results to its product portfolio in three key areas: • Rapid and scalable service provisioning: ONAPP initially intended to push the rapid provisioning to its Federation and CDN products. With realignments in the company strategy there has been more emphasis in using SUPERFLUIDITY technology within its MicroVisor and Enterprise portfolio. The MicroVisor was initially developed as part of the FP7 EUROSERVER project but work on optimisations has been carried out in SUPERFLUIDITY. The light-weight hypervisor technology of the MicroVisor is ideally suited for running and deploying RFBs and micro-services as promoted in SUPERFLUIDITY. Platform improvements have been made to improve the scalability and the number of micro-services that can be supported; • Energy reduction was initially an expected output from SUPERFLUIDITY but as work was carried out the plan has changed to look at improving the consolidation of services on a server by increasing the scale it can cope with. Energy reduction therefore is no longer a direct output but is an indirect benefit. The emphasis therefore will be on marketing the improved performance and scale enabled by SUPERFLUIDITY. • Improvements to the OnApp User Interface / dashboard to cope with and manage the expected scale of virtualised functions. Although this was initially seen as a small piece of work, as the project has evolved it has become clear that being able to visualise the scale of resources envisaged in 5G needs improvements in the UI. To this effect a completely new UI that offers an alternative to OpenStack Horizon is being developed. Some of this work has been shown in the context of T6.1 work items. Coupled with the visualisation aspects are the orchestration facilities to be able to manage the large number of services and adjust the configuration to allow better pinning and affinities to NUMA cores to allow for better performance. • Work on container support is another new large area that the SUPERFLUIDITY project has allowed ONAPP to branch out into. Although not fully integrated with OpenStack, ONAPP has moved to supporting OpenStack with its MicroVisor work. The research has led to increased experience in Kubernetes / Mesos / Docker Swarm and other container management engines that will ultimately be tied up into ONAPP product range and allow for the management of far larger numbers of end-user VMs than was possible before. The main motivation for OnApp to exploit SUPERFLUIDITY results is the belief that moving services towards the end-user is a key trend that is visible from the Cloud and Datacenter perspective. OnApp has a huge client base (3,000+ clouds, 900+ service providers in 93 countries) that use the OnApp CDN platform to expand their infrastructure beyond their own physical hardware, to bring content closer to their end-users. In addition, the OnApp Federation (that was a result of the FP7 Trilogy2 project where

	OnApp worked together with Intel, NEC and BT, http://trilogy2.eu/) is being used to allow customers to provision computation near to end-users. Adoption of the Federation and CDN platform highlight the requirement from customers to bring services close to customers. The current system requires a-priori provisioning of full Virtual Machines. By reducing the size of Virtual Machines as is being carried out in SUPERFLUIDITY, is allowing for more rapid, dynamic provisioning. Work on container support, the visualisation aspects and the underlying platform improvements have allowed for an order of magnitude more services to be managed than before. The optimisation of the Hypervisor to allow higher performance between guests has enabled more fluid use of resources available throughout the infrastructure. This combined with lighter VMs and containers has allowed for better dynamicity of services to respond, rapidly to end-user demands and reduce the amount of over-provisioning of resources currently needed to be able to respond to demand.
ALB	ALB's exploitation plan focus in transferring SUPERFLUIDITY project outputs to Altice Labs' Business Units (BU), enhancing the ALB's products and services portfolio in the long run. In particular, ALB is deeply involved in the Multi-acess Edge Computing (MEC) and Network Functions Virtualization (NFV) technologies, following and contributing to the ongoing ETSI standardization work. Within the scope of the SUPERFLUIDITY project, ALB developed a MEC prototype with respective orchestration, in order to validate the concept and show Edge Computing benefits. This prototype is the basis for a Proof-of-Concept (PoC) being conducted at the main operator (PT/MEO), in a close to production environment, demonstrating the technology and identifying opportunities, allows ALB to acquire know-how and eventually establish the foundations for a future ALB MEC solution. As the MEC technology is emerging, ALB believes that it can take a significant advantage regarding other competitors, as today full edge computing solutions are not available in the market. The relationship of MEC with other important technologies like SDN and NFV (e.g. CRAN, EPC) is also an area where ALB acquired know-how, validating and prototyping this integration. The SUPERFLUIDITY project deployed CRAN and MEC as virtualized entities, showing that together, ultra-low latency values and bandwidth efficiencies, can be achieve at a reasonable cost. Furthermore, both technologies can share the cloud resources, taking advantage of management synergies and reducing overall costs. The SUPERFLUIDITY project implemented the MEC architecture from scratch. It intends to be able to fully manage and orchestrate the whole MEC solution by the end of the project. In addition to the ETSI MEC

	architecture, the SUPERFLUIDITY implementation has the following key properties: • Compatibility with existing mobile networks - there is no need for network changes; the network is MEC-unaware. • Compatibility with existing applications - applications don't need to change to provide services, unless they want to take advantage of MEC Services (APIs). • Transparency to end users - the user is not aware about whether the service is provided from the core or the edge (they just feel different latencies). • Integration with NFV infrastructure - the operator can take advantage of NFV and MEC synergies, reducing overall costs. *MEC TOF reused by SELFNET: One of the key aspects of SELFNET is the capability to deploy network functions, either sensors and/or actuators, physical and/or virtual, at several segments of the network path (access, aggregation, core, etc.). For SELFNET, one of the most important segments to deploy network functions is at the mobile network edge, guaranteeing that access to these functions from the mobile terminals does not have any performance issues. To achieve this aim, it is required to de-encapsulate and encapsulate the mobile network packets, in particular the Generic Tunnelling Protocol (GTP). This will guarantee that network packets are de-encapsulated before being processed by the deployed network functions at the mobile edge, and encapsulated again before being redirected to the destination. Since SUPERFLUIDITY has designed and developed a software component for de-encapsulating and encapsulating GTP packets, named Traffic Offloading Function (TOF), SELFNET has decided to reuse it for the implementation of the SELFNET use-cases - Self-Optimization, Self-Healing and Self-Protection. Currently, the TOF software component is already integrated in the Altice Labs SELFNET testbed in Aveiro, Portugal.*
REDHAT	*REDHAT is planning to use SUPERFLUIDITY results to improve its "Red Hat Enterprise Linux OpenStack Platform" (see https://access.redhat.com/products/red-hat-enterprise-linux-openstack-platform), in particular the Kuryr project to enable containers to run on the same OpenStack networks as OpenStack VMs.* *Such platform delivers an integrated foundation to create, deploy, and scale a secure and reliable public or private OpenStack cloud. Red Hat Enterprise Linux OpenStack Platform combines the world's leading enterprise Linux and the fastest-growing cloud infrastructure platform to give customers the ability to scale and quickly meet customer demands without compromising on availability, security, or performance.* *SUPERFLUIDITY results will enhance the capabilities of such platform, currently focused on cloud developments, with the possibility to use and*

	deploy virtualised functions (be it as VMs, or as containers) across different appliances located in the wireless access domain next to e.g., base stations.
TELCA	TELCARIA is a young SME dealing with network service virtualization for network operators. SUPERFLUIDITY technologic achievements will allow exploit the new ecosystem created by SUPERFLUIDITY by developing and bringing to the market innovative and ground-breaking services and applications that can take advantage of the information on radio network capabilities and conditions available at the base station. This will enable TELCARIA to do rapid prototyping, design validation and direct deployment on its customers at a much faster pace than traditional vendors thus enabling gaining of market share and creation of new jobs. The two main areas in which TELCARIA foresees exploitation potential of SUPERFLUIDITY: • Advances in technologies for integrating heterogeneous wireless networks (up to RAN level) and in architectures to optimise the reuse of functionality across heterogeneous access technologies for 5G. In this context, Telcaria will strengthen future collaborations and contracts as new services in relation with 5G. • *Northbound APIs for user management and network control. The APIs defined by SUPERFLUIDITY for services instantiation, as well as tools such as RDCL 3D, will help easy deployments for network operators. Moreover, we will likely integrate these APIs in Alviu, a Telcaria's product focused on this area and based on the ONOS framework.*
TID	*During the lifetime of the project, Telefonica has continued the deployment of its UNICA infrastructure, already available in the most relevant of our operations worldwide. In parallel, the company has heavily engaged in the development of an open-source production ready NFV orchestration framework, OSM (https://osm.etsi.org/) hosted by ETSI, and that is about to become the orchestration platform for the UNICA infrastructure.* *Some results from SUPERFLUIDITY are being incorporated into OSM, and TID will continue contributing to the community code base aligned with work initiated in the project, as it is the case of the translation of NEMO-based intent declarations into NFV descriptors and orchestration scripts. In addition, the evolution of UNICA has made Telefonica start considering the deployment of small-footprint datacentres and the use of whiteboxes as forwarding elements. TID is in the process of applying results from SUPERFLUIDITY in these two aspects, in particular in what relates to the characterization and benchmarking of these infrastructure elements.*
USTR	SUPERFLUIDITY will provide USTR with a unique opportunity to improve the streaming/muxing software within CDN setups at the edge of the network. The targeted product will be in particular the UNIFIED REMIX, USTR product that creates personalised streams based on a recommendation, a ruleset or even an external repository such as an ad

	network. The opportunity of using SUPERFLUIDITY's results is to mux content as late as possible has with the benefit of reducing latency and lower the amount of CDN internal traffic thus reducing cost at operations level. USTR will bring its A/V streaming expertise to the project as well as access to the professional streaming solutions it creates and provide use-case testing and integration. More specific exploitation plans include: Location awareness created by virtualised Mobile Edge Computing support in SUPERFLUIDITY: • improvement of streaming/muxing software in the edge (scripts to deploy unified streaming software in the cloud) • remix, personalization, enabling costumers better opportunities for monetization • transmuxing in realistic virtualised edge computing, reducing backend traffic and caching requirements • better reception/video optimization based on Radio Access information • these will all reduce network transmission and caching cost for customers running our software The scale independence offered by the Nokia CloudBand and the Intel telemetry setup: • enable a video streaming origin that runs in a private cloud and scales compute/network resources on the go automatically • This will benefit Unified Streaming customers that currently typically runs software on public clouds or private servers NEC work on fast instantiation and minimalistic kernels: • Will enable unified streaming to run on as minimalistic kernels as possible, enabling USTR to offer highly cost efficient implementations of unified origin in a cloud environment. • Availability in the cloud SUPERFLUIDITY provides USTR a key opportunity to deploy their video streaming software based on state of the art virtualization and cloud technologies. Experiments with unified streaming using Docker, unikernel, Xen, KVM, rancher and other cloud software will enable unified streaming to offer its video streaming services in different environments. This offers new opportunities for Revenues
ACADEMIC PARTNERS	Academic partners plan the exploitation of SUPERFLUIDITY results in order to further increase their standings and ranking in the academic area through participation to key scientific conferences. Furthermore, research results from SUPERFLUIDITY will flow into the academic curricula by endorsing solutions and new paradigms in teaching programmes. Both MSc. and PhD students will be targeted. Teaching material will be updated with the latest technological achievements so as to improve the quality of the offered studies. Last but not least, academic partners will use

	SUPERFLUIDITY's results for attracting industry consulting contracts and to launch new innovative software-based academic spin-offs. As a matter of fact, members of CNIT involved in SUPERFLUIDITY just founded a spin-off, together with some talented, newly graduated, engineers. The spin-off is already selling mobile applications developed for the health sector and it is developing applications for the cultural heritage and tourism industry. The participation to SUPERFLUIDITY could offer significant opportunities to the spin-off, which focuses on the development of applications using cutting edge mobile technologies. For example, the capability of doing efficient real time streaming and sharing of multimedia information on mobile devices will be factored in future products of the spin-off.

Table 3: Per-partner Exploitation Plans

3.2 Consortium as a Whole

The overall exploitation plan of the consortium as a whole is reported in Table 4 below.

EXPLOITATION PLAN: CONSORTIUM AS A WHOLE

In terms of consortium, SUPERFLUIDITY is bringing partners with complementary expertise addressing different roles in the exploitation of project results, here we present the overall vision of the consortium starting from bottom up (from hardware, through operating system up to services and their orchestration including the architectural aspects):
• INTEL: standard high volume elements (client, edge and server platforms, networking equipment); *supporting software capabilities such as telemetry (e.g. snap), packet processing acceleration (DPDK);*
• NEC: integration of software solutions for small cells, LTE base stations and cloud resources;
• NOKIA
  o *NOKIA FR: Integration of Nokia Xhaul MPR product in 5G market; Enhance the list of supported features in the fronthaul catalogue of Nokia (re-programmability, SDN based control); Enhance the Cloud RAN V2.0 product with technology as container, Unikernel, Ansible, Kubernetes; Support of Service based architecture like microservice and RFB in the Next Generation CORE;*
  o NOKIA IL: management tools for orchestrating virtualised functions across heterogeneous access;
• REDHAT: *operating systems, OpenStack, OpenShift and orchestration running on top, such as ManageIQ and Ansible;*
• TELCARIA: integrate heterogeneous wireless networks information for running innovative services;
• CITRIX: virtualised *load balancing and* traffic management, spreading over the converged infrastructure (up to the access);

• ONAPP: virtualised storage services spreading over the converged infrastructure (up to the access). *ONAPP has also gained experience in working with INTEL and UNIFIED STREAMING on the telemetry, analytics and working with video streaming content*; • UNIFIED STREAMING: streaming services spreading over the converged infrastructure (up to the access); • TELECOM OPERATORS (TID, ALB, BT): converged infrastructure moving service functionalities between core and access.

Table 4: Exploitation Plan of the Consortium as a Whole

3.3 Proof of Concept Activities

To demonstrate innovation in practice and prove feasibility of the exploitation plans above, partners have executed a multitude of Proof of Concept (PoC) activities, as outlined in Table 5 below. Many of these PoC activities were executed during the second half of the project, particularly as part of WP7.

PARTNER	POC DEMOS/PLANS	STATUS
NOKIA FR	• C-RAN and packet core deployment based on Docker containers • Intent based fronthaul solution for C-RAN • Automation of the deployment of end to end mobile network: RAN and CORE using RDCL developed in WP6 • *Integration of Nokia xHaul MPR product in the fronthaul link* • *Re-programmability of the fronthaul link via Intent based networking*	Demonstrated to top 5 tier-one customers. Demonstrated in Future X Days 2016[1]. Demonstrated in 5G event Rome. *Demonstrated in EUCNC 2017.* *Demonstrated in Future X Days 2017.* *Demonstrated to Mobile Network Business Unit 2018.*
ALB	• Prototype Multi-access Edge Computing (MEC) platform	*Installed and demonstrated to PT/MEO Network Strategy Unit, at their own 4G/WLAN laboratory, integrated with Huawei and*

[1] Open Nokia Bell-Labs event for Customers and Business Unit (Villarceaux, October 4, 2016)

		Ericsson eNBs. The Demo is based on project components, aligned with WP7 use case/scene 3.b. This PoC was recently updated and the latest MEC prototype version, allowing multiple access technologies and applications chaining. In that context an AR application from a Portuguese third-party was demonstrated, using both 4G and WLAN access networks. Results are documented in D7.3.
USTR	• Late Transmuxing in Mobile Edge Computing (MEC)	Demonstrated as part of 1st review.
CITRIX	• Virtualization/cloudification of EPC/GiLAN applications, interoperable with the major cloud orchestration options and aligned with the *MEC and* NFV architecture • *Load balancing and* subscriber-aware service chaining, employing NSH (Network Service Header) for scalable and efficient delivery of composite network services	*TCP optimizer as VNF, leveraging NetScaler, was integrated with the project MEC platform and validated in WP7.* *OpenStack-integrated load balancer, using NetScaler LBaaS provider, was also validated in WP7.*
ONAPP	• Will trial the low-resource Hypervisors (MicroVisor) within the ONAPP platform to investigate how performance can be improved against most recent product version • ONAPP has created a new UI and developed support in its Cloud product to make it work with OpenStack. This significant effort allows for managing far greater number of services and for working with heterogeneous resources as the MicroVisor works with both ARMv8 and Intel x86 hardware architectures. • ONAPP is also involved in the testbed development and overall project PoC planning and integration efforts.	The combination of the ONAPP platform with OpenStack is a critical part of its new, planned Enterprise range and therefore is well known to senior stakeholders.

TELCA	• Intends to participate in the testing and evaluation of SUPERFLUIDITY results with a PoC on virtualization, aligned with the NFV framework of ETSI. In addition, thanks to their cooperation with REDIRIS (the National Research Network, NREN, in Spain), it is possible to test the PoC with captures of real traffic, which will help to make a more realistic activity. • Plans to leverage the orchestration framework of SUPERFLUIDITY and the different NBI developed, for an internal product called Alviu.	To be installed in the 5TONIC lab, as part of the collaborations with UC3M and IMDEA.
CNIT	• *Is involving Ninux.org, an experimentation oriented wireless community network in Italy, in the evaluation on the field of the Superfluidity concepts, exemplified in the A-RFB PoC, described in D7.3.*	*Tested in CNIT lab. Just released to Ninux.org at the time of writing.*
UPB	• Prepared a SymNet based verification demo that currently works on top of real networks. Working on verifying networks specified with OpenStack Neutron in a 10-node rack + switch.	Demonstrated as part of 1st review.
INTEL	• Automated KPI Mapping to support service assurance based optimisation of Telco KPI's. • Implementation of an automated methodology, which can support reasoning over the trade-offs between total cost of ownership (TCO), service performance and the available infrastructure resource landscape (Block Abstractions). • *Automated Fingerprint for workload class determination.* • *Utility functions for workload on-boarding optimisation.*	*Demonstrated to Business Unit, PoC activity completed with Business Unit.* *Tech transfer to BU completed.*
REDHAT	• *Hybrid deployment of OpenShift on top of Baremetal Servers and OpenStack VMs, with kuryr enabled. Based on OpenStack TripleO Installer and triggered from ManageIQ*	*Presented at DevConf 2018.*
NOKIA IL	• *Proof of concept of the Integration of the newly developed optimal placement procedure into Nokia's NFVO.*	*Demonstrated to internal partners.*

Table 5: Proof of Concept Activities

4 Conclusion

As reported by this deliverable, the SUPERFLUIDITY project has successfully achieved its objectives, in terms of executing the plan of activities in the areas of innovation and commercial exploitation.

To support the above statement, below are some measurable innovation and exploitation outcomes:

- Research bibliography (state of the art) at project inception: ~85 items (see 5 References)
- # of research publications advancing the state of the art: ~45 (see Section 2.3 & 5 References)
- # of citations to the above research publications: ~185 (albeit many of the publications are new)
- # of exploitation actions throughout the project duration: 15 (see Reports for Periods 4 to 8)
- # of proof-of-concept (PoC) activities demonstrated: 16 (more than planned, a few still pending)

Judging from the above statistics, we can conclude that:

- Superfluidity innovations, as published at major research conferences, workshops and journals, are already being swiftly diffused across the scientific community, fueling additional research and innovation activities.
- Project outcomes have been actively promoted within Superfluidity member organizations, particularly ones with a commercial orientation, reaching important stakeholders and decision makers. As indicated by the exploitation actions and PoC status updates, there is great interest towards adopting Superfluidity outcomes in products and solutions of the respective partners.

Based on the above observations, we are confident that Superfluidity results will have a lasting impact, embedded in commercial offerings and driving further innovations, post project completion.

5 References

[8021Q] IEEE Std. 802.1Qbg-2012, "IEEE Standard for Local and Metropolitan Area Networks — Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks—Amendment 21: Edge Virtual Bridging," July 5, 2012, http://standards.ieee.org/getieee802/download/802.1Qbg-2012.pdf, p. 191.

[AIN2016] Khartik Ainala, Rufael N. Mekuria, Birendra Khathariya, Zhu Li, Ye-Kui Wang, Rajan Joshi: "An improved enhancement layer for octree based point cloud compression with plane projection approximation", SPIE Applications of Digital Image Processing XXXIX, 99710R (September 27, 2016).

[ARA2017] G. Aravinthan, H. Dalia, C.S. Chen, L. Roullet: "Virtualization of Radio Access Network by Virtual Machine and Docker: Practice and Performance Analysis", IFIP/IEEE Integrated Network Management Symposium (IM 2017), May 8-12, 2017, Lisbon.

[ARN2010] Applied Technology and Innovation Management, Heinrich Arnold, Michael Erner, Peter Möckel, Christopher Schläffer Editors, Springer Heidelberg Dordrecht London New York

[ASG2016] H. J. Asghar, L. Melis, C. Soldani, E. De Cristofaro, M. A. Kaafar, and L. Mathy: "SplitBox: Toward Efficient Private Network Function Virtualization" ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization (HotMiddlebox 2016), Co-located with ACM SIGCOMM'16, August 22-26, 2016, Salvador Florianópolis, Brazil.

[ATT2016] AT&T "ECOMP (Enhanced Control, Orchestration, Management & Policy) Achitecture" White Paper, March 2016

[BAN12] M. Bansal, J. Mehlman, S. Katti, and P. Levis, "OpenRadio: A Programmable Wireless Dataplane," in Proc. of HotSDN'12, Aug. 2012.

[BAN15] Manu Bansal, Aaron Schulman, and Sachin Katti. 2015. Atomix: a framework for deploying signal processing applications on wireless infrastructure. In Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation (NSDI'15). USENIX Association, Berkeley, CA, USA, 173-188.

[BEL2010] Beloglazov, Anton, and Rajkumar Buyya. "Energy efficient resource management in virtualized cloud data centers." 10th IEEE/ACM Int. Conf. on Cluster, Cloud and Grid Computing. IEEE Computer Society, 2010.

[BHA2013] Bhatnagar, Vaibhav, et al. "An FPGA Software Defined Radio Platform with a High-Level Synthesis Design Flow." Vehicular Technology Conference (VTC Spring), 2013 IEEE 77th. IEEE, 2013.

[BHAN10] G. Bhanage, I. Seskar, R. Mahindra, and D. Raychaudhuri, "Virtual Basestation: Architecture for an Open Shared WiMAX Framework," in Proc. of VISA'10, 2010.

[BIA12] D. Garlisi, F. Giuliano, I. Tinnirello, P. Gallo, F. Gringoli and G. Bianchi, "Deploying virtual MAC protocols over a shared access infrastructure using MAClets," Computer Communications Workshops (INFOCOM WKSHPS), 2013 IEEE Conference on, Turin, 2013, pp. 17-18.

[BIA14] Giuseppe Bianchi, Marco Bonola, Antonio Capone, Carmelo Cascone: OpenState: programming platform-independent stateful openflow applications inside the switch. ACM Comp. Commun. Rev. 44(2): 44-51 (2014)

[BiaArx14] Bianchi, G., Bonola, M., Picierro, G., Pontarelli, S., & Monaci, M. (2013). StreaMon: a data-plane programming abstraction for Software-defined Stream Monitoring. arXiv preprint arXiv:1311.2442.

[BON2017a] Marco Bonola, Roberto Bifulco, Luca Petrucci, Salvatore Pontarelli, Angelo Tulumello, Giuseppe Bianchi: "Implementing advanced network functions for datacenters with stateful programmable data planes", IEEE LANMAN 2017 – The 23rd IEEE International Symposium on Local and Metropolitan Area Networks, Osaka, Japan, June 12-14, 2017.

[BON2017b] Marco Bonola, Roberto Bifulco, Luca Petrucci, Salvatore Pontarelli, Angelo Tulumello, Giuseppe Bianchi: "Demo: implementing advanced network functions with stateful programmable data planes", IEEE LANMAN 2017 – The 23rd IEEE International Symposium on Local and Metropolitan Area Networks, Osaka, Japan, June 12-14, 2017.

[BOS13] P. Bosshart, G. Gibb, H.-S. Kim, G. Varghese, N. McKeown, M. Izzard, F. Mujica, and M. Horowitz, "Forwarding metamorphosis: Fast programmable match-action processing in hardware for sdn," in ACM SIGCOMM 2013. ACM, 2013, pp. 99–110.

[BOT2012] Botero, Juan Felipe, et al. "Energy efficient virtual network embedding." Communications Letters, IEEE 16.5 (2012): 756-759.

[BRU2018] Marco Spaziani Brunella, Salvatore Pontarelli, Marco Bonola, Giuseppe Bianchi: "V-PMP: a VLIW Packet Manipulator Processor", EuCNC2018 – NET. Ljubljana, Slovenia, June 18-21, 2018.

[CHE15] A. Checko et al., "Cloud RAN for Mobile Networks—A Technology Overview," in IEEE Communications Surveys & Tutorials, vol. 17, no. 1, pp. 405-426, Firstquarter 2015.

[CHI2016] Luca Chiaraviglio, Nicola Blefari-Melazzi, William Liu, Jairo A. Gutierrez, Jaap Van De Beek, Robert Birke, Lydia Chen, Filip Idzikowski, Daniel Kilper, Paolo Monti, Jinsong Wu: "5G in Rural and Low-Income Areas: Are We Ready?", ITU Kaleidoscope Conference, Bangkok, Thailand, November 2016.

[CHI2017a] Luca Chiaraviglio, Nicola Blefari-Melazzi, William Liu, Jairo A. Gutierrez, Jaap van de Beek, Robert Birke, Lydia Chen, Filip Idzikowski, Daniel Kilper, Paolo Monti, Antoine Bagula, Jinsong Wu. "Bringing 5G in Rural and Low-Income Areas: Is it Feasible?" IEEE Communications Standards Magazine, 2017.

[CHI2017b] Luca Chiaraviglio, Nicola Blefari Melazzi, Carla-Fabiana Chiasserini, Bogdan Iatco, Francesco Malandrino, Stefano Salsano: "An Economic Analysis of 5G Superfluid Networks", 18th IEEE International Conference on High Performance Switching and Routing (IEEE HPSR), Campinas, Brazil, June 18-21, 2017.

[CHI2017c] Luca Chiaraviglio, Lavinia Amorosi, Stefania Cartolano, Nicola Blefari Melazzi, Paolo Dell'Olmo, Mohammad Shojafar, Stefano Salsano: "Optimal Superfluid Management of 5G Networks", IEEE NetSoft 2017, Bologna, Italy, 3-7 July 2017.

[CHI2018] Luca Chiaraviglio, Fabio D'Andreagiovanni, Giulio Sidoretti, Nicola Blefari-Melazzi, Stefano Salsano: "Optimal Design of 5G Superfluid Networks: Problem Formulation and Solutions", ICIN 2018 – 21st International Conference on Innovation in Clouds, Internet and Networks, Paris, France, February 2018.

[COH2016] Alejandro Cohen, Asaf Cohen and Omer Gurewitz, "Secure Group Testing", IEEE International Symposium on Information Theory (ISIT), Barcelona, 10-15, July 2016.

[Cra13] B. Mack-Crane. "OpenFlow Extensions". In: US Ignite ONF GENI workshop, October 8, 2013.

[DAR2017] Tooska Dargahi, Alberto Caponi, Moreno Ambrosin, Giuseppe Bianchi, Mauro Conti. "A Survey on the Security of Stateful SDN Data Planes". IEEE Communications Surveys and Tutorials, 30 March 2017.

[DET2013] Gregory Detal, Benjamin Hesmans, Olivier Bonaventure, Yves Vanaubel, and Benoit Donnet. 2013. Revealing middlebox interference with tracebox. In Proceedings of the 2013 conference on Internet measurement conference (IMC '13).

[DET2016] Andrea Detti, Bruno Ricci, Nicola Blefari-Melazzi: "Tracker-assisted rate adaptation for MPEG DASH live streaming". Proceedings of the IEEE International Conference on Computer Communications (INFOCOM) 2016. 2016, April 10-15. San Francisco, United States of America.

[DOB2009] M. Dobrescu, N. Egi, K. Argyraki, B.-G. Chun, K. Fall, G. Iannaccone, A. Knies, M. Manesh and S. Ratnasamy, "RouteBricks: Exploiting Parallelism to Scale Software Routers", ACM SOSP'09, 2009.

[DON2012] Dong, Yaozu, et al. "High performance network virtualization with SR-IOV." Journal of Parallel and Distributed Computing 72.11 (2012): 1471-1480.

[DPDK2016] "DPDK: Data Plane Development Kit – What it is". dpdk.org. Retrieved June 2016.

[DUN09] R. Duncan and P. Jungck, "PacketC: Language for High-Performance Packet Processing", IEEE HPCC'09, 2009.

[ESP2017] Flavio Esposito, Andrej Cvetkovski, Tooska Dargahi, and Jianli Pan: "Complete Edge Function Onloading for Effective Backend-driven Cyber Foraging", IEEE International Conference on Wireless and Mobile Computing, Networking and Communications – WiMob 2017, Rome, Italy, October 9-11, 2017.

[ETT2016] G. Bianchi, E. Biton, N. Blefari-Melazzi, I. Borges, L Chiaraviglio, Pedro de la Cruz Ramos, P. Eardley, F. Fontes, M. J. McGrath, L. Natarianni, D. Niculescu, C. Parada, M. Popovici, V. Riccobene, S. Salsano, B. Sayadi, J. Thomson, C. Tselios, G. Tsolis, "SUPERFLUIDITY: A Flexible Functional Architecture for 5G Networks", Transactions on Emerging Telecommunications Technologies, Wiley.

[FAWG] Openflow Forwarding Abstractions Working Group Charter. http: //goo.gl/TtLtw0. Apr. 2013.

[FOS11] Foster, Nate, et al. "Frenetic: A network programming language." ACM SIGPLAN Notices. Vol. 46. No. 9. ACM, 2011.

[GON2017a] Roberto Gonzalez, Filipe Manco, Alberto Garcia-Duran, Jose Mendes: “Net2Vec: Deep Learning for the Network”, ACM SIGCOMM Workshop on Big Data Analytics and Machine Learning for Data Communication Networks, Big-DAMA 2017, Los Angeles, USA, August 20-25, 2017.

[GON2017b] Roberto Gonzalez, Alberto Garcia-Duran, Filipe Manco, Mathias Niepert: “Network Data Monetization Using Net2Vec“, ACM SIGCOMM 2017, Los Angeles, USA, August 20-25, 2017.

[GUD2013] Aditya Gudipati, Daniel Perry, Li Erran Li, Sachin Katti, “SoftRAN: Software Defined Radio Access Network.” ACM HotSDN’13, August 16, 2013, Hong Kong, China.

[HAN2010] S. Han, K. Jang, K. Park, and S. Moon, “PacketShader: a GPU-Accelerated Software Router”, ACM SIGCOMM'10, 2010.

[HON2011] Michio Honda, Yoshifumi Nishida, Costin Raiciu, Adam Greenhalgh, Mark Handley, and Hideyuki Tokuda. 2011. Is it still possible to extend TCP?. In Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference (IMC '11).

[HUA14] C. L. I, J. Huang, R. Duan, C. Cui, J. Jiang and L. Li, "Recent Progress on C-RAN Centralization and Cloudification," in IEEE Access, vol. 2, no. , pp. 1030-1039, 2014.

[HWA2015] J. Hwang, K. K. Ramakrishnan and T. Wood, "NetVM: High Performance and Flexible Networking Using Virtualization on Commodity Platforms," in IEEE Transactions on Network and Service Management, vol. 12, no. 1, pp. 34-47, March 2015.

[JEY13] V. Jeyakumar, M. Alizadeh, C. Kim, and D. Mazieres, “Tiny packet programs for low-latency network control and monitoring,” in ACM Workshop on Hot Topics in Networks (HOTNETS 2013), 2013.

[JEY2013] Vimalkumar Jeyakumar, Mohammad Alizadeh, David Mazières, Balaji Prabhakar, Changhoon Kim, and Albert Greenberg. 2013. EyeQ: practical network performance isolation at the edge. In Proceedings of the 10th USENIX conference on Networked Systems Design and Implementation

[KAM2016] Joseph Kampeas, Asaf Cohen and Omer Gurewitz, "On Secrecy Rates and Outage in Multi-User Multi-Eavesdroppers MISO Systems", IEEE International Symposium on Information Theory (ISIT), Barcelona, July 10-15, 2016.

[KOH2000] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, “The Click Modular Router”, ACM Transactions on Computer Systems 18(3), Aug 2000.

[KOK10] R. Kokku, R. Mahindra, H. Zhang, and S. Rangarajan, “NVS: a virtualization substrate for WiMAX networks,” in Proc. of MobiCom’10, 2010.

[LAC2013] Katrina LaCurts, Shuo Deng, Ameesh Goyal, and Hari Balakrishnan. 2013. Choreo: network-aware task placement for cloud applications. In Proceedings of the 2013 conference on Internet measurement conference (IMC '13).

[LUN2018] Paolo Lungaroni, Claudio Pisa, Stefano Salsano, Giuseppe Siracusano, Francesco Lombardo: "Extending OpenVIM R3 to support Unikernels (and Xen)", ETSI – Open Source Mano Conference. Rome, Italy. February 06 - 09, 2018.

[MAD2013] Unikernels: Library operating systems for the cloud. A. Madhavapeddy, R. Mortier, C. Rotsos, D. Scott, B. Singh, T. Gazagnaire, S. Smith, S. Hand, and J. Crowcroft. SIGPLAN Not., 48(4):461–472, Mar. 2013.

[MAE11] Maeder, A., Mancuso, V., Weizman, Y., Biton, E., Rost, P., Perez-Costa, X., & Gurewitz, O., "FLAVIA: Towards a generic MAC for 4G mobile cellular networks", Future Network & Mobile Summit, 2011

[MAI2012] S. Maital, D. V. R. Seshadri, Innovation Management, Response Books

[MAN2017] F. Manco, C. Lupu, F. Schmidt, J. Mendes, Simon Kuenzer, S. Sati, K. Yasukata, C. Raiciu, F. Huici: "My VM is Lighter (and Safer) than your Container", ACM SOSP 17 – 26th Symposium on Operating Systems Principles, Shangai, China, 28-31 October 2017.

[MAR2014] ClickOS and the Art of Network Function Virtualization. J. Martins, M. Ahmed, C. Raiciu, V. Olteanu, M. Honda, R. Bifulco, F. Huici. NSDI 2014.

[MEC] "Mobile Edge Computing – Introductory Technical White Paper". https://portal.etsi.org/Portals/0/TBpages/MEC/Docs/Mobile-edge_Computing_-_Introductory_Technical_White_Paper_V1%2018-09-14.pdf

[MEK2016a] Rufael Mekuria, Michael McGrath, Christos Tselios, Dirk Griffioen, George Tsolis, Shahar Beiser: "KPI Mapping for Virtual Infrastructure Scaling for a Realistic Video Streaming Service Deployment", 8th International Conference on Quality of Multimedia Experience (IEEE QoMEX), Lisbon, Portugal, June 6-8, 2016.

[MEK2016b] R. Mekuria, J. Fennema, R. Belleman, D. Griffioen, USTR – Unified Streaming: "Multi-Protocol Video Delivery with Late Trans-Muxing" ACM Multimedia, 15 - 19 October 2016, Amsterdam, The Netherlands.

[MEK2016c] Rufael Mekuria, Pablo Cesar: "MP3DG-PCC, Open Source Software Framework for Implementation and Evaluation of Point Cloud Compression", ACM Multimedia Conference 2016, 15-19 October 2016, Amsterdam, The Netherlands.

[MEK2017] R. Mekuria, S. Laserre and C. Tulvan: "Performance assessment of point cloud compression", IEEE VCIP 2017 – Visual Communication and Image Processing, St. Petersburg, FL. December 10 – 13 , 2017.

[MEK2018] R. Mekuria, M. J. McGrath, V. Bayon-Molino, V. Riccobene, C. Tselios, A. Dobrodub, J. Thomson: accepted paper "Automated Profiling of Virtualized Media Processing Functions using Telemetry and Machine Learning", ACM MMSys 2018 - ACM Multimedia Systems Conference. Amsterdam, the Netherlands. June 12 - 15, 2018.

[Mey13] D. Meyer, "Openflow: Today's reality, tomorrow's promise? an architectural perspective," in SDN Summit 2013, Paris, France, March 2013, 2013.

[MIJ2016a] R. Mijumbi, J. Serrat, J. L. Gorricho, N. Bouten, F. De Turck and R. Boutaba, "Network Function Virtualization: State-of-the-Art and Research Challenges," in IEEE Communications Surveys & Tutorials, vol. 18, no. 1, pp. 236-262, Firstquarter 2016.

[MIJ2016b] R. Mijumbi, J. Serrat, J. l. Gorricho, S. Latre, M. Charalambides and D. Lopez, "Management and orchestration challenges in network functions virtualization," in IEEE Communications Magazine, vol. 54, no. 1, pp. 98-105, January 2016.

[MON13] Monsanto, Christopher, et al. "Composing Software Defined Networks." NSDI. 2013.

[MOR99] Robert Morris, Eddie Kohler, John Jannotti, and M. Frans Kaashoek. 1999. The Click modular router. SIGOPS Oper. Syst. Rev. 33, 5 (December 1999), 217-231. DOI=10.1145/319344.319166 http://doi.acm.org/10.1145/319344.319166

[MOS14] M. Moshref, A. Bhargava, A. Gupta, M. Yu, and R. Govindan. Flow-level State Transition as a New Switch Primitive for SDN. In Proceedings of the ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN'14), August 2014.

[MUE2010] Mueck, Markus, et al. "ETSI reconfigurable radio systems: status and future directions on software defined radio and cognitive radio standards." Communications Magazine, IEEE 48.9 (2010): 78-86.

[MWH2014] I. Marinos, R. N. M. Watson, and M. Handley. Network Stack Specialization for Performance. In Proceedings of the ACM SIGCOMM 2014 Conference, pages 175–186, 2014.

[NFV001] ETSI GS NFV-SWA 001 V1.1.1 "Network Functions Virtualisation (NFV); Virtual Network Functions Architecture", December 2014

[NFV007] ETSI GS, "Network Functions Virtualisation (NFV); Infrastructure; Methodology to describe Interfaces and Abstractions"; NFV-INF 007 V1.1.1 (2014-10)

[NFV2014] ETSI Network Functions Virtualisation, http://www.etsi.org/technologies-clusters/technologies/nfv

[NGMN15] NGMN Alliance "Further Study on Critical C-RAN Technologies", White Paper, March 31, 2015

[NSM2016] "Hybrid IP/SDN networking: open implementation and experiment management tools", S. Salsano, P. L. Ventre, F. Lombardo, G. Siracusano, M. Gerola, E. Salvadori, M. Santuari, M. Campanella, L. Prete. IEEE Transactions on Network and Service Management.

[NVGRE] M. Sridharan et al., "NVGRE: Network Virtualization Using Generic Routing Encapsulation," IETF Draft- sridharan-virtualization-nvgre-03.txt, Aug. 2013

[NVO] Narten et al., "Problem Statement: Overlays for Network Virtualization," IETF Draft draft-ietf-nvo3-overlay-problem-statement-04, July 31, 2013

[ODPL] OpenDataPlane. http://www.opendataplane.org

[OEC] OpenEdgeComputing, available online: http://openedgecomputing.org/index.html

[OF1.4] Open Networking Foundation "Openflow specification, ver 1.4", October 2013

[OSF2016] OpenStack Foundation "Accelerating NFV Delivery with Openstack", Technical Report, February 2016

[PAL2015] Shoumik Palkar, Chang Lan, Sangjin Han, Keon Jang, Aurojit Panda, Sylvia Ratnasamy, Luigi Rizzo, and Scott Shenker. 2015. E2: a framework for NFV applications. In Proceedings of the 25th Symposium on Operating Systems Principles (SOSP '15). ACM, New York, NY, USA, 121-136. DOI=http://dx.doi.org/10.1145/2815400.2815423

[PAR2018] Carlos Parada, Francisco Fontes, Carlos Marques, Vitor Cunha, Cristina Leitão: “Multi-access Edge Computing: A 5G Technology”, EuCNC2018 – NET. Ljubljana, Slovenia, June 18-21, 2018.

[PEN13] K. Pentikousis, Y. Wang, and W. Hu, “Mobileflow: Toward software defined mobile networks,” Communications Magazine, IEEE, vol. 51, no. 7, pp. 44–53, 2013.

[Per13] P. Peresini, M. Kuzniar, and D. Kostic. “OpenFlow Needs You! A Call for a Discussion About a Cleaner OpenFlowAPI”. In Proc. of the EU Workshop on Software Defined Network (EWSDN). 2013.

[PICSIG] PCI-SIG, “Single Root I/O Virtualization and Sharing 1.1 Specification,” http://www.pcisig.com/members/down-loads/specifications/iov/sr-iov1_1_20Jan10.pdf, available only to members.

[PON2017] Salvatore Pontarelli, Marco Bonola, Giuseppe Bianchi: “Smashing SDN “built-in” actions: programmable data plane packet manipulation in hardware”, IEEE NetSoft 2017, Bologna, Italy, July 3-7, 2017.

[POP2010] Lucian Popa, Minlan Yu, Steven Y. Ko, Sylvia Ratnasamy, and Ion Stoica. 2010. CloudPolice: taking access control out of the network. In Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks (Hotnets-IX)

[POP2012] Lucian Popa, Gautam Kumar, Mosharaf Chowdhury, Arvind Krishnamurthy, Sylvia Ratnasamy, and Ion Stoica. FairCloud: sharing the network in cloud computing. In Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication (SIGCOMM '12)

[POP2013] Lucian Popa, Praveen Yalagandula, Sujata Banerjee, Jeffrey C. Mogul, Yoshio Turner, and Jose Renato Santos. 2013. ElasticSwitch: practical work-conserving bandwidth guarantees for cloud computing. SIGCOMM Comput. Commun. Rev. 43

[POP2016] Matei Popovici, Costin Raiciu: “Exploiting Multipath Congestion Control for Fun and Profit”, HOTNETS 2016, Atlanta, Georgia

[PRE2017a] Ioannis Prevezanos, Andreas Angelou, Christos Tselios, Alexandros Stergiakis, Vassilis Tsogkas and George Tsolis: “Hammer: A Real-world, end-to-end Network Traffic Simulator”, IEEE CAMAD 2017, 22nd IEEE International Workshop on Computer Aided Modelling and Design of Communication Links and Networks, Lund, Sweden, 19-21 June 2017.

[PRE2017b] I. Prevezanos, C. Tselios, A. Angelou, M. McGrath, R. Mekuria, V. Tsogkas and G. Tsolis: “Evaluating Hammer Network Traffic Simulator: System Benchmarking and Testbed Integration”, IEEE GLOBECOM 2017 – Global Hub: Connecting East and West, Singapore, 4-8 December 2017.

[QAZ2013] Zafar Ayyub Qazi, Cheng-Chun Tu, Luis Chiang, Rui Miao, Vyas Sekar, and Minlan Yu. 2013. SIMPLE-fying middlebox policy enforcement using SDN. SIGCOMM Comput. Commun. Rev. 43, 4 (August 2013)

[RAD15] Radisys Corp. “Evaluating Cloud RAN Implementation Scenarios” White Paper, June 2015

[RAI2012] Costin Raiciu, Mihail Ionescu, and Dragoş Niculescu. 2012. Opening up black box networks with CloudTalk. In Proceedings of the 4th USENIX conference on Hot Topics in Cloud Computing (HotCloud'12).

[RAM2009] Ram, Kaushik Kumar, et al. "Achieving 10 Gb/s using safe and transparent network interface virtualization." Proc. 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments. ACM, 2009.

[RFC6325] R. Perlman et al., "Routing Bridges (RBridges): Base Protocol Specification," IEEE RFC 6325, July 2011, 99 pages, http://tools.ietf.org/html/

[RIZ2012] L. Rizzo, "Netmap: a novel framework for fast packet i/o. In Proceedings of the 2012 USENIX conference on Annual Technical Conference", USENIX ATC'12, USENIX Association, Berkeley, CA, USA, 2012.

[RIZ2012a] L. Rizzo, M. Carbone, G. Catalli, "Transparent acceleration of software packet forwarding using netmap", in Proceedongs of IEEE INFOCOM'12, 2012.

[ROJ2018] Elisa Rojas, Roberto Doriguzzi-Corin, Sergio Tamurejo, Andres Beato, Arne Schwabe, Kevin Phemius, Carmen Guerrero. "Are We Ready to Drive Software-Defined Networks? A Comprehensive Survey on Management Tools and Techniques". ACM Computing Surveys (CSUR), March 2018.

[SACH08] J. Sachs and S. Baucke, "Virtual radio: a framework for configurable radio networks," in Proc. of WICON'08, Nov. 2008.

[SAL2017a] Stefano Salsano, Francesco Lombardo, Claudio Pisa, Pierluigi Greto, Nicola Blefari Melazzi, "RDCL 3D, a Model Agnostic Web Framework for the Design and Composition of NFV Services", 3rd IEEE International Workshop on Orchestration for Software Defined Infrastructures, O4SDI at IEEE NFV-SDN conference, Berlin, 6-8 November 2017, arXiv:1702.08242v21.

[SAL2017b] S. Salsano, L. Chiaraviglio, N. Blefari-Melazzi, C. Parada, F. Fontes, R. Mekuria, D. Griffioen: "Toward Superfluid Deployment of Virtual Functions: Exploiting Mobile Edge Computing for Video Streaming", Soft5 Workshop, 1st International Workshop on Softwarized Infrastructures for 5G and Fog Computing, in conjunction with 29th ITC conference, Genoa, Italy, September 8, 2017. Presentation.

[SAT15] Satyanarayanan, M., Schuster, R., Ebling, M., Fettweis, G., Flinck, J., Joshi, K., Sabnani, K. "An Open Ecosystem for Mobile-Cloud Convergence", in IEEE Communications Magazine, Volume 53, Number 3, March 2015

[SCH2016] Arne Schwabe, Pedro A. Aranda Gutiérrez, Holger Karl: "Composition of SDN applications: Options/challenges for real implementations", Applied Networking Research Workshop 2016 (ANRW'16), 17-22 July 2016, Berlin, Germany.

[SCH2017] Florian Schmid: "uniprof: A Unikernel Stack Profiler", ACM SIGCOMM 2017, Los Angeles, USA, August 20-25, 2017.

[SEK2012] Design and implementation of a consolidated middlebox architecture. V. Sekar, N. Egi, S. Ratnasamy, M. K. Reiter, and G. Shi. 9th USENIX conf. on Networked Systems Design and Implementation, NSDI 2012.

[SHE2012b] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, V. Sekar, "Making middleboxes someone else's problem: network processing as a cloud service", SIGCOMM, 2012.

[SHO2017] Mohammad Shojafar, Luca Chiaraviglio, Nicola Blefari-Melazzi, Stefano Salsano: "P5G: A Bio-inspired Algorithm for the Superfluid Management of 5G Networks", IEEE GLOBECOM 2017 – Global Hub: Connecting East and West, Singapore, 4-8 December 2017.

[SIR2016] Giuseppe Siracusano, Roberto Bifulco, Simon Kuenzer, Stefano Salsano, Nicola Blefari Melazzi, Felipe Huici: "On-the-Fly TCP Acceleration with Miniproxy", ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization (HotMiddlebox 2016), Co-located with ACM SIGCOMM'16, August 22-26, 2016, Salvador Florianópolis, Brazil.

[SIR2017] Giuseppe Siracusano, Roberto Bifulco, Stefano Salsano: "TCP Proxy Bypass: All the Gain with No Pain!", ACM SIGCOMM 2017, Los Angeles, USA, August 20-25, 2017.

[SON13] H. Song, "Protocol-oblivious forwarding: Unleash the power of sdn through a future-proof forwarding plane," 2nd ACM SIGCOMM HotSDN Workshop, 2013, pp. 127–132.

[STO2016a] Radu Stoenescu, Matei Popovici, Lorina Negreanu, and Costin Raiciu, UPB – University Politehnica of Bucharest: "SymNet – Scalable Symbolic Execution for Modern Networks", ACM SIGCOMM 2016, Florianópolis, Brazil, August 22-26, 2016.

[STO2016b] Radu Stoenescu, Dragos Dumitrescu, Costin Raiciu: "OpenStack Networking for Humans: Symbolic Execution to the Rescue" invited paper, IEEE LANMAN 2016, Rome, Italy, June 13 -15, 2016.

[STT] B. Davie, Ed., J. Gross, "A Stateless Transport Tunneling Protocol for Network Virtualization (STT)," IETF Draft draft-davie-stt-03.txt, Mar. 12, 2013, 19 pages, http://tools.ietf.org/html/draft-davie-stt-03

[TAJ2018] Mohammad Mahdi Tajiki, Stefano Salsano, Mohammad Shojafar, Luca Chiaraviglio, Behzad Akbari: "Energy-efficient Path Allocation Heuristic for Service Function Chaining", ICIN 2018 – 21st International Conference on Innovation in Clouds, Internet and Networks, Paris, France, February 2018.

[TAN11] K. Tan et al., "Sora: High-Performance Software Radio Using General-Purpose Multi-Core Processors," Communications of the ACM, vol. 54, pp. 99–107, Jan. 2011.

[TIN12] I. Tinnirello, G. Bianchi, P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, "Wireless MAC Processors: Programming MAC Protocols on Commodity Hardware", IEEE INFOCOM, March 2012.

[TSE2016] C. Tselios and G. Tsolis "On QoE-awareness Using Virtualized Probes in 5G Networks", IEEE CAMAD 2016, October 23-25, 2016, Toronto, Canada.

[VAS2017] Laura Vasilescu, Vladimir Olteanu, Costin Raiciu: "Sharing CPUs via endpoint congestion control", IEEE SIGCOMM KBNets 2017, Los Angeles, US-CA, August 21-25, 2017.

[VEI2017] Paul Veitch, Edel Curley and Tomasz Kantecki: “Performance Evaluation of Cache Allocation Technology for NFV Noisy Neighbour Mitigation”, IEEE NetSoft 2017, Bologna, Italy, 3-7 July 2017.

[VEN2016] Pier Luigi Ventre, Claudio Pisa, Stefano Salsano, Giuseppe Siracusano, Florian Schmidtz, Paolo Lungaroni, Nicola Blefari Melazzi: “Performance Evaluation and Tuning of Virtual Infrastructure Managers for (Micro) Virtual Network Functions”, 2016 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), 7-9 November 2016, Palo Alto, California, USA.

[VXLAN] M. Mahalingam et al. “VXLAN: A Framework for Over-laying Virtualized Layer 2 Networks over Layer 3 Net-works,” IETF Draft draft-mahalingam-dutt-dcops-vxlan-04.txt, May 8, 2013.

[VZN2016] Verizon “SDN-NFV Reference Architecture v1.0” Verizon Network Infrastructure Planning, February 2016

[WAG2017] Arjen Wagenaar, Dirk Griffioen, Rufael Mekuria: “Unified Remix: a Server Side Solution for Adaptive Bit-Rate Streaming with Inserted and Edited Media Content”, ACM Multimedia Systems Conference, Taipei, Taiwan, June 20-24, 2017

[WU2014] Wu, Wenfei, et al. "PRAN: Programmable Radio Access Networks." 13th ACM HotNet Workshop, 2014.

[Yad11] N. Yadav and D. Cohn. “OpenFlow Primitive Set”, available online at http://goo.gl/6qwbg. July 2011.

[YAP10] K.-K. Yap et al., “Blueprint for introducing innovation into wireless mobile networks,” Proc. of VISA’10, 2010.

[YAS2017] K. Yasukata, F. Huici, V. Maffione, G. Lettieri, M. Honda: “HyperNF: Building a High Performance, High Utilization and Fair NFV Platform”, SoCC ’17 – ACM Symposium on Cloud Computing 2017, Santa Clara, CA, 25-27 September 2017.

[ZAK10] Y. Zaki, L. Zhao, C. Grg, A. Timm-Giel, “A Novel LTE Wireless Virtualization Framework,” MONAMI’10.

[ZEN2012] Hongyi Zeng, Peyman Kazemian, George Varghese, and Nick McKeown. 2012. Automatic test packet generation. In Proceedings of the 8th international conference on Emerging networking experiments and technologies (CoNEXT '12).