

Leveraging In-Network Computing with Network Function Virtualization

Elias P. Duarte Jr. (Federal University of Parana)

Curitiba, Brazil

www.inf.ufpr.br/elias

elias@inf.ufpr.br



Keynote presented at

**17th ACS/IEEE International Conference
on Computer Systems and Applications
(AICCSA 2020) Nov. 5th, 2020**

Agenda


- Network Softwarization: Virtualization
- Network Function Virtualization: NFV
- In-Network-Computing
- INC using Programmable Hardware
- *NFVinc*: Advantages
- Case studies: NFV-FD, VNF-Consensus, NFV-RBCast
- Final Remarks

Network Virtualization

- A true revolution: network equipment and resources usually available in hardware can be implemented in software
 - *For instance*: instead of having a single physical router → commodity hardware can run multiple different virtual routers
- Virtual networks are created, provisioned, managed and terminated in software
 - Multiple different virtual networks can co-exist on the same physical infrastructure

Network Virtualization

Virtual Networks A, B and C are on the same physical infrastructure but can run completely different protocol stacks and are mutually isolated from a security point of view



Internet Ossification: Solved! ✓

- The Internet currently connects nearly 5 billion human beings around the world (total pop. ~7.8bi)
→ <https://www.internetworldstats.com/stats.htm>
- Protocols developed decades ago were threatening the evolution of the Internet
- With virtualization: the network can evolve to multiple directions at the same time, all of which can co-exist

Not only the Internet!

- Virtualization is changing all networks, not only the Internet




Virtualization: Enabling Technologies

- 1) Software Defined Networking (SDN)**
- 2) Data Plane Programmability**
- 3) Network Function Virtualization (NFV)**

Let's check what they are about:


Software Defined Networking (SDN)

- The heart of SDN technology lies on *flow-level models of the Internet*
- A flow is a sequence of packets, usually sharing the same origin/destination (IP addresses and ports) and protocol



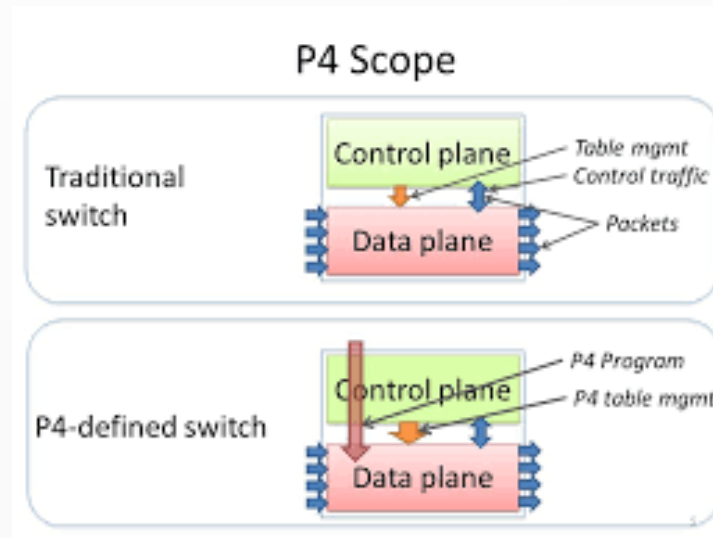
SDN: OpenFlow

- OpenFlow is *the* current SDN enabling technology
- The most important feature: **separation** of the **control plane** from the **data plane**
- Data Plane: switches that simply forward data
- Control Plane: runs all control tasks



Data Plane Programmability

- Switches on the data plane usually have a predefined, static set of instructions
- What if you can make the set of instructions dynamic?
- Program the instructions a switch will execute
- Most prevalent enabling technology: P4




Network Function Virtualization

- The root of NFV technology is on *middleboxes*
- Middleboxes are also called “network appliances” or “network functions”
- They are intermediate devices on the path from an IP source and an IP destination
- Run a myriad of services: IDS, firewall, NAT, proxies, caches,..



NFV: Virtual Network Functions

- NFV was first proposed to allow the implementation of middleboxes in software, using virtualization technology
- Instead of specialized/often proprietary hardware: run virtual functions on off-the-shelf hardware
- Reduces costs: OPEX (Operational Expenses) & CAPEX (Capital Expenses)
- Improves flexibility, management, saves energy




A Deep Market Change

- A shift from buying specialized hardware from a handful of vendors to downloading services from Internet Marketplaces
- <http://coral.ufsm.br/gt-fende/en/home.html>



SFC: Service Function Chain

- Besides Individual VNFs (Virtual Network Functions)
- Complex network services can be constructed called Service Function Chains by orchestrating multiple VNFs



NFV-MANO: Standards for NFV Interoperability

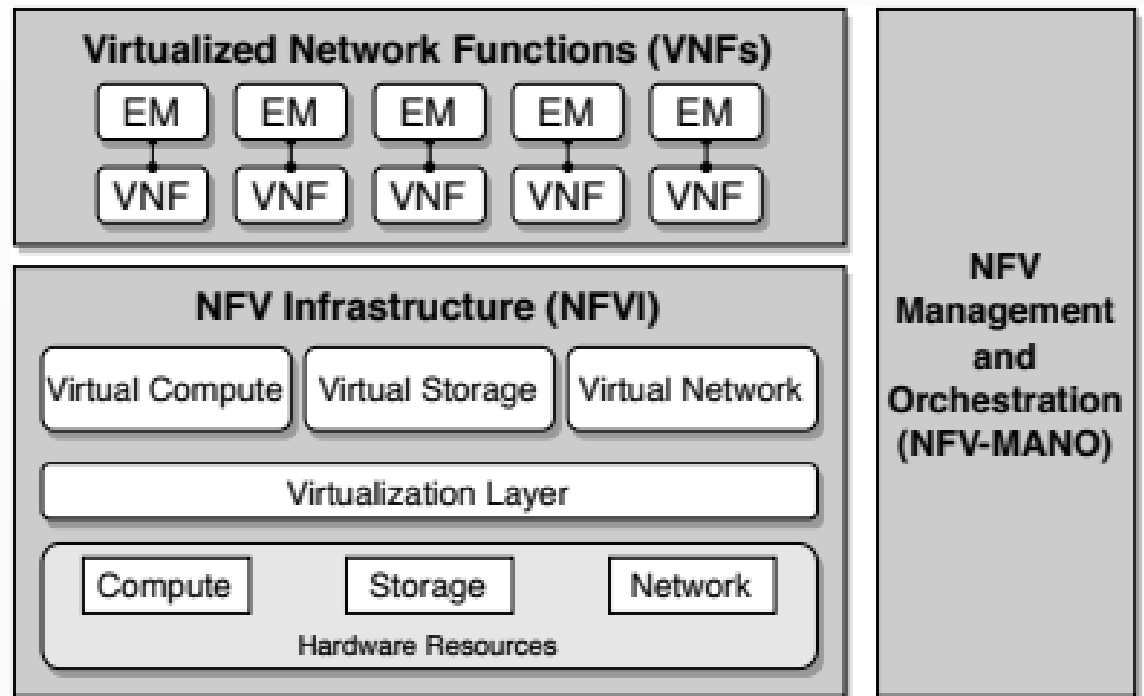
- The ETSI has proposed an NFV reference model:
NFV-MANO: NFV Management and Orchestration
- Objective: standardize the execution and management of network services based on NFV
- Providing interoperability for VNFs and SFCs from different vendors (actually providers)

Several NFV Platforms Based on NFV-MANO

- Multiple open source platforms are available that are based on NFV-MANO:
 - OSM (Open Source MANO)
 - OpenBaton
 - OPNFV (OPen NFV)
 - Tacker (OpenStack project)
 - Vines (for CloudStack)
<http://www.inf.ufpr.br/jwvflauzino/vines/>

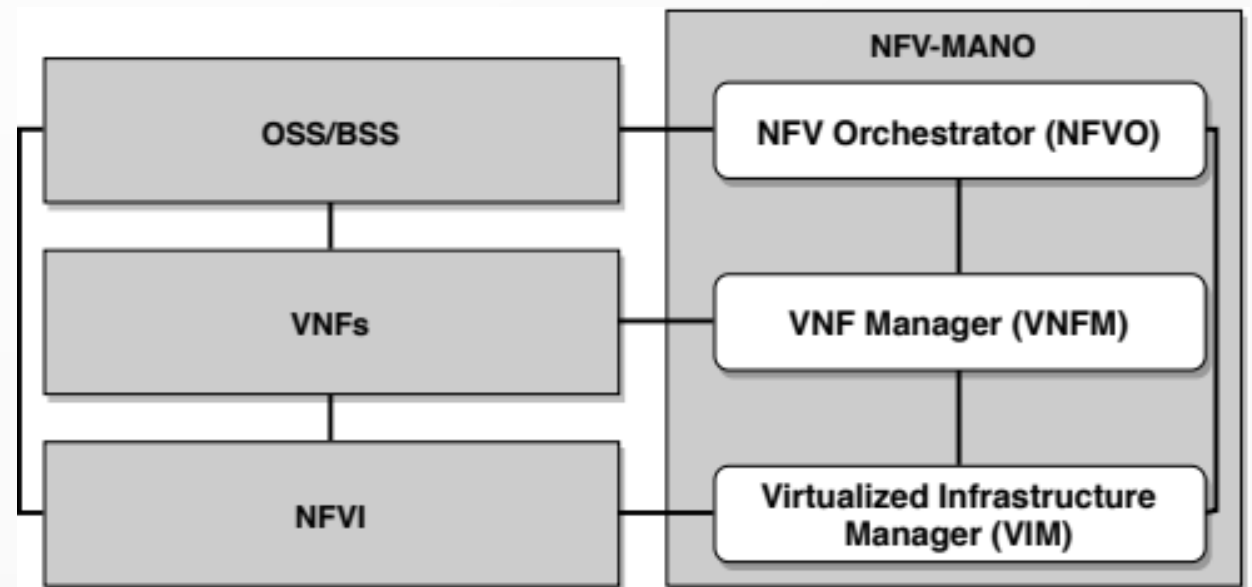
NFV-MANO: Architecture

- The reference model is composed of 3 main blocks: NFV-MANO (itself), the NFVI, and the VNFs
- NFVI (NFV Infrastructure) usually is the virtual infrastructure on which VNFs are
 - instantiated
 - executed
 - managed



The NFV-MANO Block

- The NFV-MANO block consists of modules for VNF orchestration and management
- **NFVM**: NFV-Manager is responsible for VNF lifecycle (VNF instantiation, configuration, update, deletion, scale in/out, monitoring, failure detection)
- **NFVO**: NFV-Orchestrator Network service lifecycle
- **VIM**: usually a cloud platform



Changing Subjects: INC

- INC is the acronym for In-Network Computing
- The idea is to move the application logic into the network
- New services implemented *within* the network
- These services, which are generally executed by hosts on the border...
- ... become network services!
- Natively offered by the network itself

The Two First INC Efforts

- In 2015 two groups proposed two different INC strategies independently:
- Our group proposed an implementation of failure detectors:
 - R. C. Turchetti and E. P. Duarte, “Implementation of a Failure Detector Based on Network Function Virtualization,” IEEE/IFIP DSN-W, 2015.
- Nearly simultaneously, Soule et. al. proposed NetPaxos:
 - H. T. Dang, D. Sciascia, M. Canini, F. Pedone, and R. Soulé, “Netpaxos: Consensus at Network Speed,” ACM SOSR’15-SIGCOMM, 2015.

INC Based on Programmable Hardware

- NetPaxos: the idea is to implement the Paxos consensus protocol on an SDN (OpenFlow) switch
- Two different flavors are presented:
 - An optimistic version of Paxos that does not require OpenFlow extensions
 - A full version of Paxos that does require to extend the OpenFlow protocol
- In 2016 the authors described the implementation of Paxos with P4 (“Paxos Made Switch-y”) also in 2020 (“P4xos: Consensus as a Network Service”)

Other INC Efforts on Hardware

- Technologies such as ASICs (Application Specific Integrated Circuits) and FPGA (Field Programmable Gate Array)
- Have been used to implement INC services such as
 - In-Network Caches
 - Data aggregation (MapReduce)
 - Among others...
- It has been shown that these INC services can reduce the latency and/or improve the throughput

NFV*inc*: Leveraging INC with NFV

- NFV technology allows the implementation of novel services within the network
- Brings all the advantages of having the services implemented in software
- Potential significant reduction of CAPEx and OPEx
- In comparison with the solutions implemented in hardware: not the same performance level
- Different niches!

NFV*inc*: User Point of View

- Target and Main advantage: instead of implementing or obtaining and running/maintaining a service locally, either as an application or middleware...
- ...it is readily available from the network, as a native network service
- All the user needs to do is to employ the correct API to invoke to service that is already there

NFV*inc*: NetOp Point of View

- Flexibility, Flexibility, Flexibility, Flexibility
 - Compare updating the version of a service implemented in software versus the same in hardware
- No physical appliances are involved, virtual software executed on commodity hardware
- Can be downloaded from a marketplace
- Want to make it available? Just start it!
- Want to terminate the service? Just stop it!
- Saves energy, space, management is so much easier

Case Studies

- Failure Detector

- R. Turchetti, Elias P. Duarte Jr., "NFV-FD: Implementation of a Failure Detector Using Network Virtualization Technology," International Journal of Network Management, 2017.

- Consensus

- G. Venancio, E. Camargo, R. Turchetti, E. P. Duarte Jr., "VNF-Consensus: : A Virtual Network Function for Maintaining a Consistent Distributed SDN Control Plane," International Journal of Network Management, 2020.

- Reliable and Ordered Broadcast

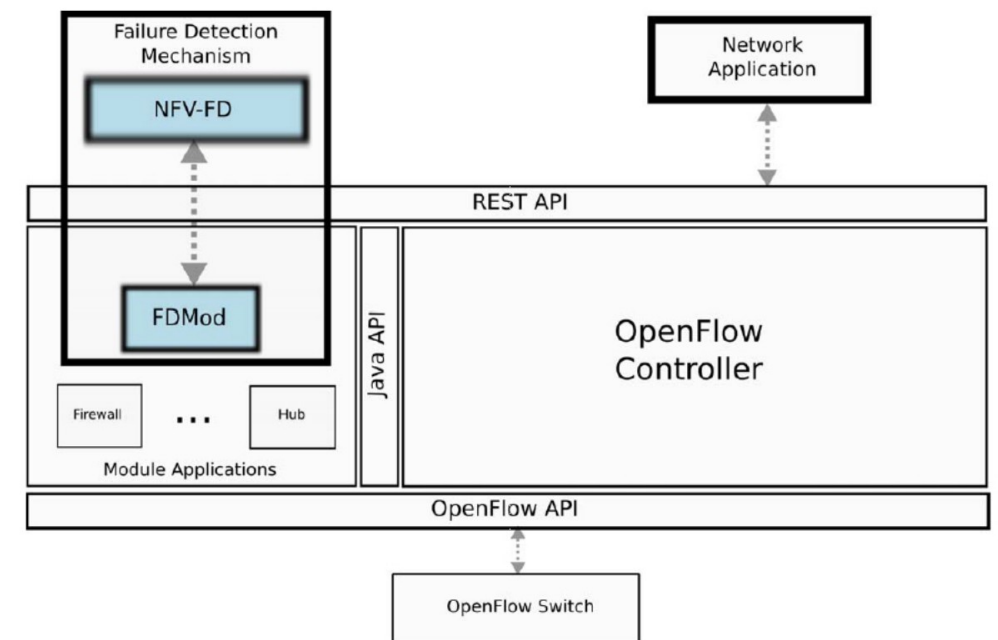
- G. Venancio, R. Turchetti, E. P. Duarte Jr., "NFV-RBCast: Enabling the Network to Offer Reliable and Ordered Broadcast Services," IEEE/SBC LADC, 2019.

Failure Detector

- A *failure detector* is a distributed systems abstraction that monitors processes
- A local module that a process accesses to obtain the states of the other processes (correct/suspected)
- Proposed in the context of the impossibility of consensus in asynchronous systems with crash faults
 - They did show that if a failure detector is employed and presents certain properties, then consensus becomes possible!

NFV-FD

- A failure detector implemented as an NFV service that makes use of information from an SDN network
- Interacts with the OpenFlow Controller
- Implicitly detects failures
- Also informs link state

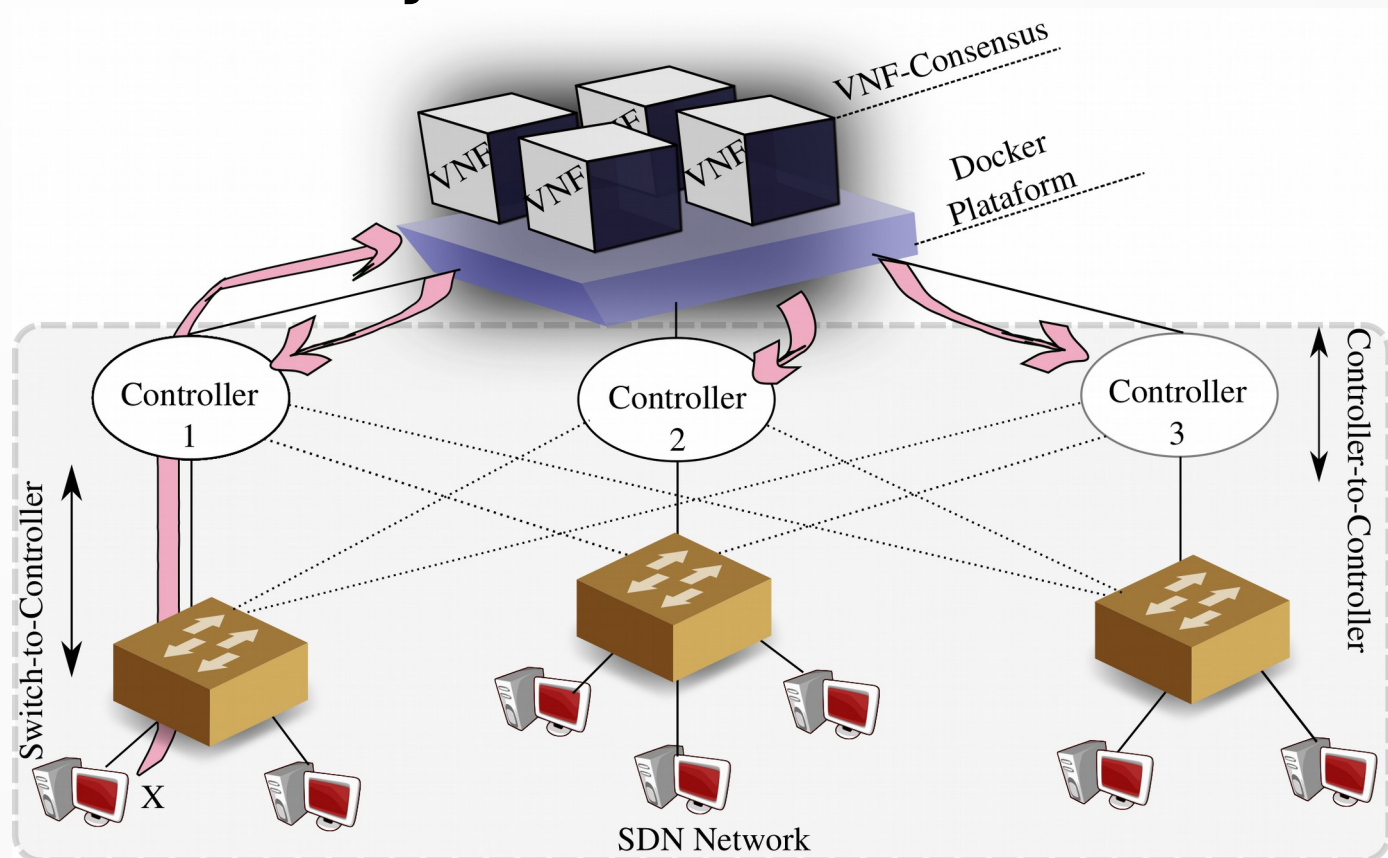


Consensus

- Allows a set of processes to reach an agreement on a value, given an initial set of possibilities
- VERY important in distributed systems: the basis for distributed replication, distributed consistency, etc.
- Paxos is one of the most important consensus algorithms
- Properties: safety (always) & liveness (under weak synchrony assumptions)

VNF-Consensus

- Proposed in the context of a distributed control plane: multiple SDN controllers
- Used to guarantee the consistency
- *Outside* the controller

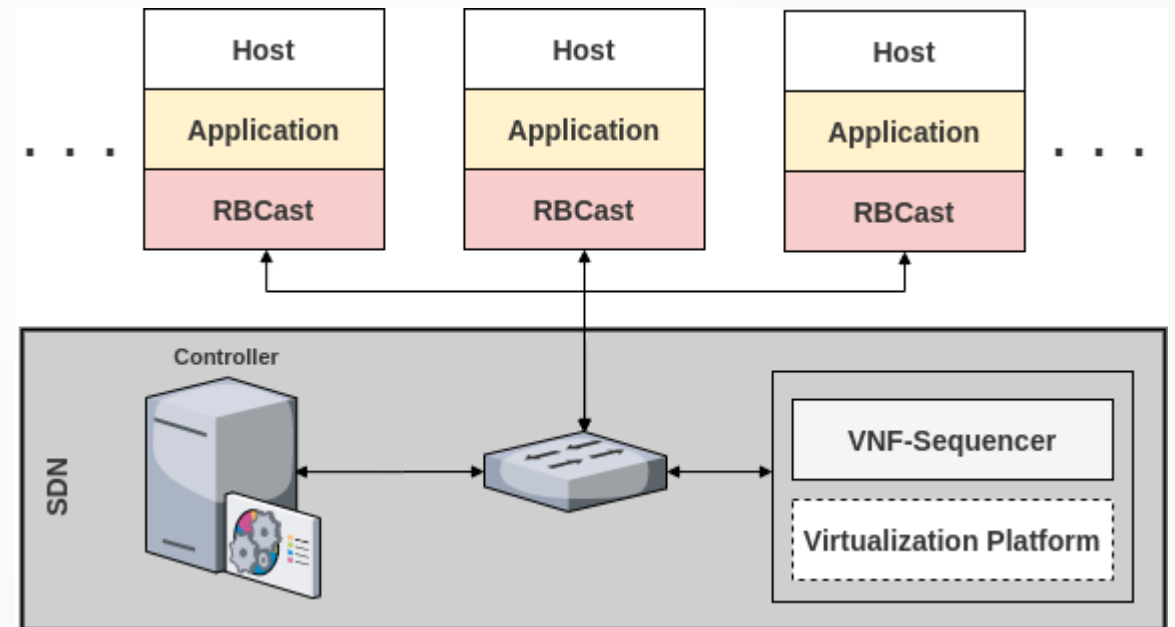


Reliable and Ordered Broadcast

- **Reliable Broadcast:** ensures that a message sent by a source process is either received by all processes or by none, despite crash faults
- **Ordered Broadcast** has another property: ensures that all processes receive all messages in a given order
 - FIFO order
 - Causal order
 - Atomic order

NFV-RBCast

- A network function that allows the network itself to offer reliable broadcast as a native service
- Also: FIFO, Causal and Atomic Broadcast
- A sequencer that is executed within the network and establishes message ordering



Conclusions

- In this work we discussed how different virtualization technologies are changing the network landscape
- Among the enabling technologies: NFV
- The synergy of NFV & INC: the network can support novel services in a flexible way
- Case studies: NFV-FD, VNF-Consensus & NFV-RBCast

Thank you!
elias@inf.ufpr.br
www.inf.ufpr.br/elias