Leveraging In-Network Computing with Network Function Virtualization

Elias P. Duarte Jr.
Universidade Federal do Paraná, Curitiba, Brazil
www.inf.ufpr.br/elias
elias@inf.ufpr.br

Research Report presented at
The 84th Meeting of IFIP W.G. 10.4
Arcos de Valdevez, Portugal
June 26, 2023

Agenda for this short presentation

- The Network Softwarization Revolution
  - Virtualization
- NFV: Network Function Virtualization
- COIN: COmputing In the Network
  - In-Network Computing: INC
- NFV+COIN = NFV-COIN → Synergy!
- Two other projects:
  - FIT-SFC: Fault- and Intrusion-Tolerant Service Function Chaining
  - HyperPaxos: Hierarchical Consensus

Network Virtualization

- Network equipment and resources available as software using virtualization
  - A true revolution, several advantages
  - *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

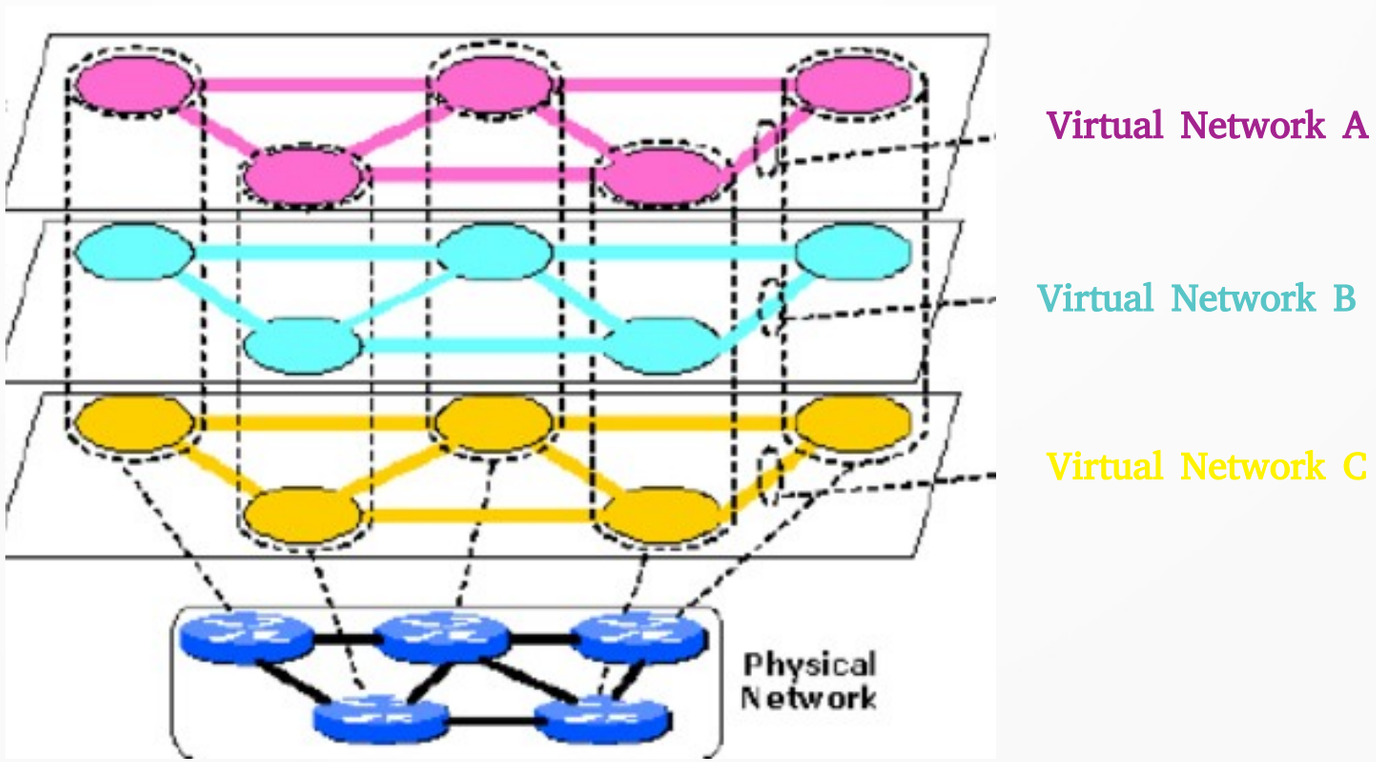

Internet Ossification: Solved! √

- The Internet currently connects around 5.5 billion human beings around the world (total pop. ~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

Network Virtualization: Enabling Technologies

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

Network Function Virtualization

- Middleboxes -- 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

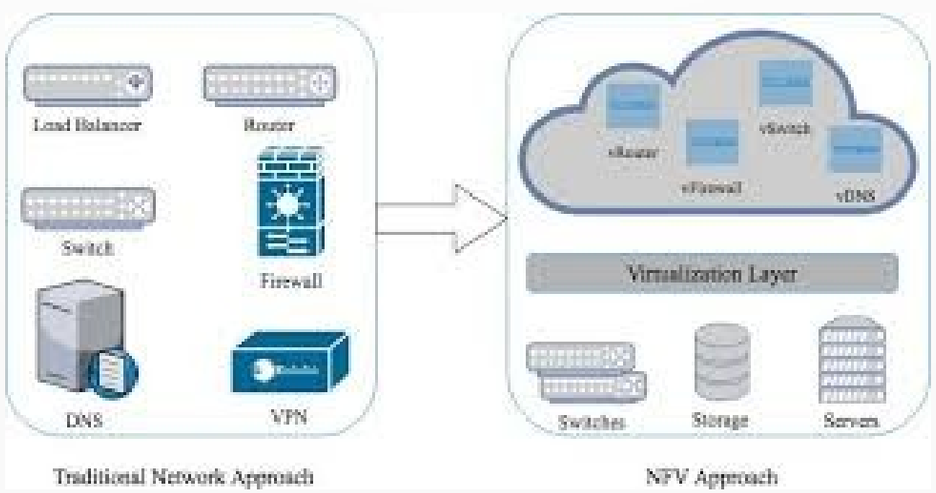

A Deep Market Change

- 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

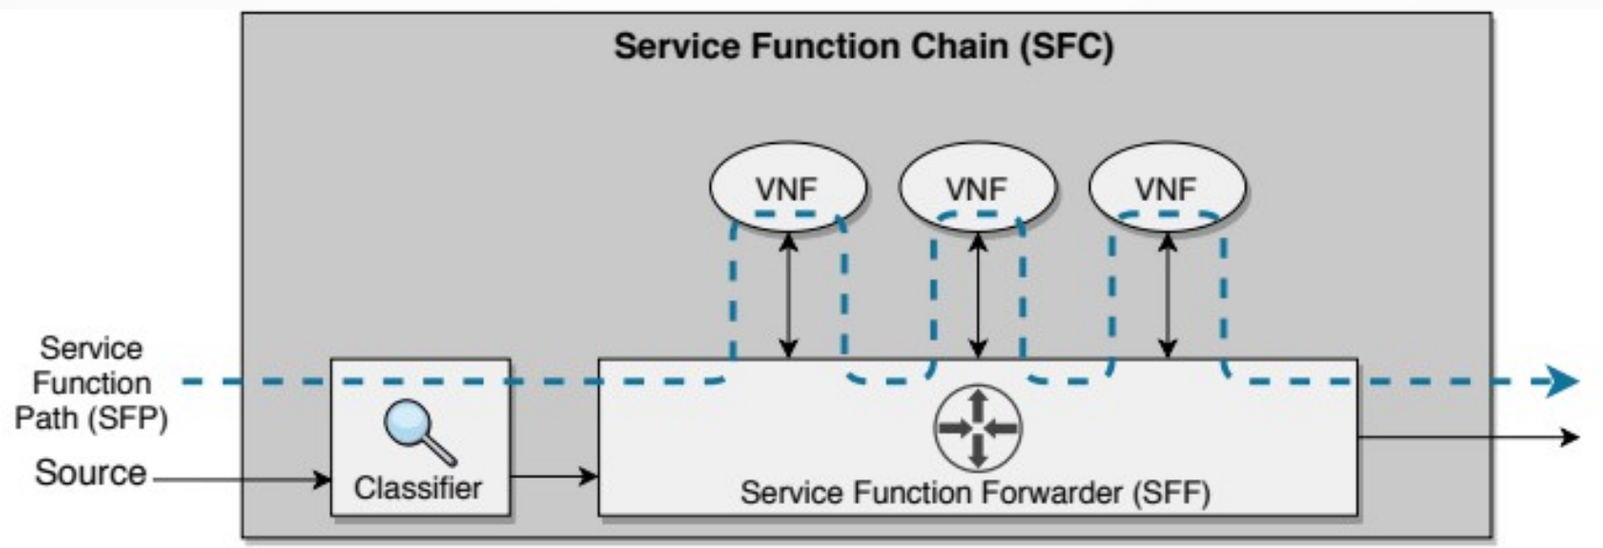

Changing Subjects: INC

- INC is the acronym for In-Network Computing
- COIN: COmputing In the Network
- 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…
- … can be 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.

Usually INC is Based 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
  - NetPaxos: the the Paxos consensus protocol on an SDN (OpenFlow) switch
  - In-Network Caches
  - In-Network Data aggregation
  - Among others...
- It has been shown that these INC services can reduce the latency and/or improve the throughput

NFV-*COIN*: Leveraging INC with NFV

- NFV technology used to deploy 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
- But a reduction of CAPEX & OPEX
- Plus: Flexibility, Flexibility, Flexibility, Flexibility

Case Studies

- Failure Detector
  - R. Turchetti, Elias P. Duarte Jr., "NFV-FD: Implementation of a Failure Detector Using Network Virtualization Technology," DSN-W, 2015
- 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.

Towards an NFV-COIN Architecture

- G. Venancio, R. Turchetti, E. P. Duarte Jr., "NFV-COIN: Unleashing the Power of In-Network Computing with NFV Technology," *Journal of Internet Services and Applications (JISA)*, Vol. 12, No. 1, pp. 46-53, 2022.

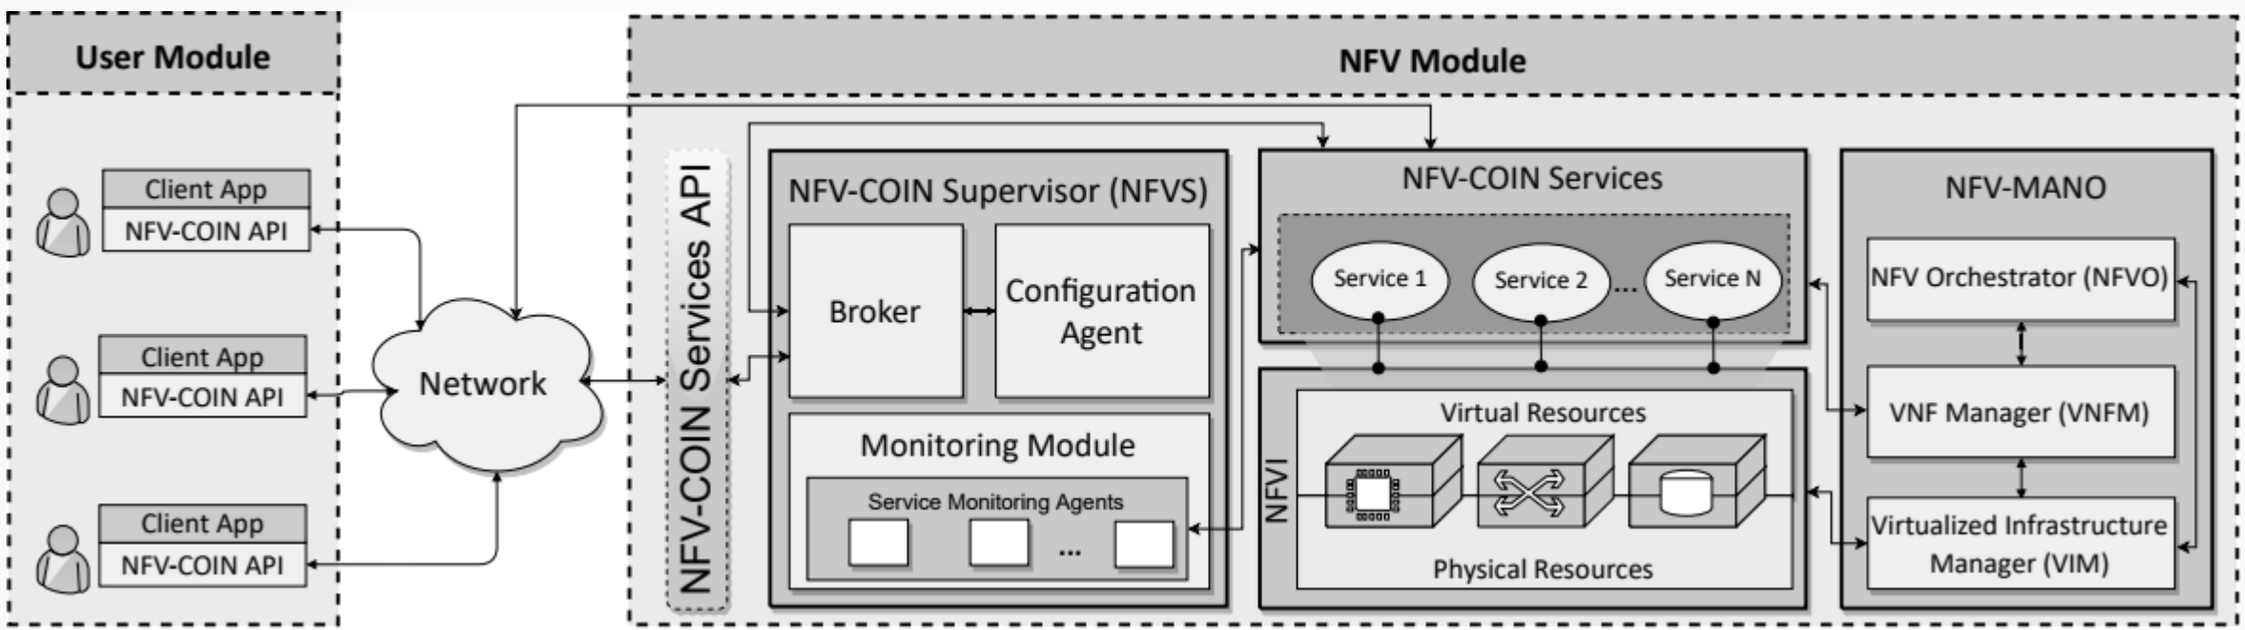

FIT-SFC

- Fault- and Intrusion Tolerant Service Function Chaining
- The IETF defined an SFC architecture
- Proposal: use replication to tolerate crash & Byzantine failures
- Taking into consideration the components of the architecture itself

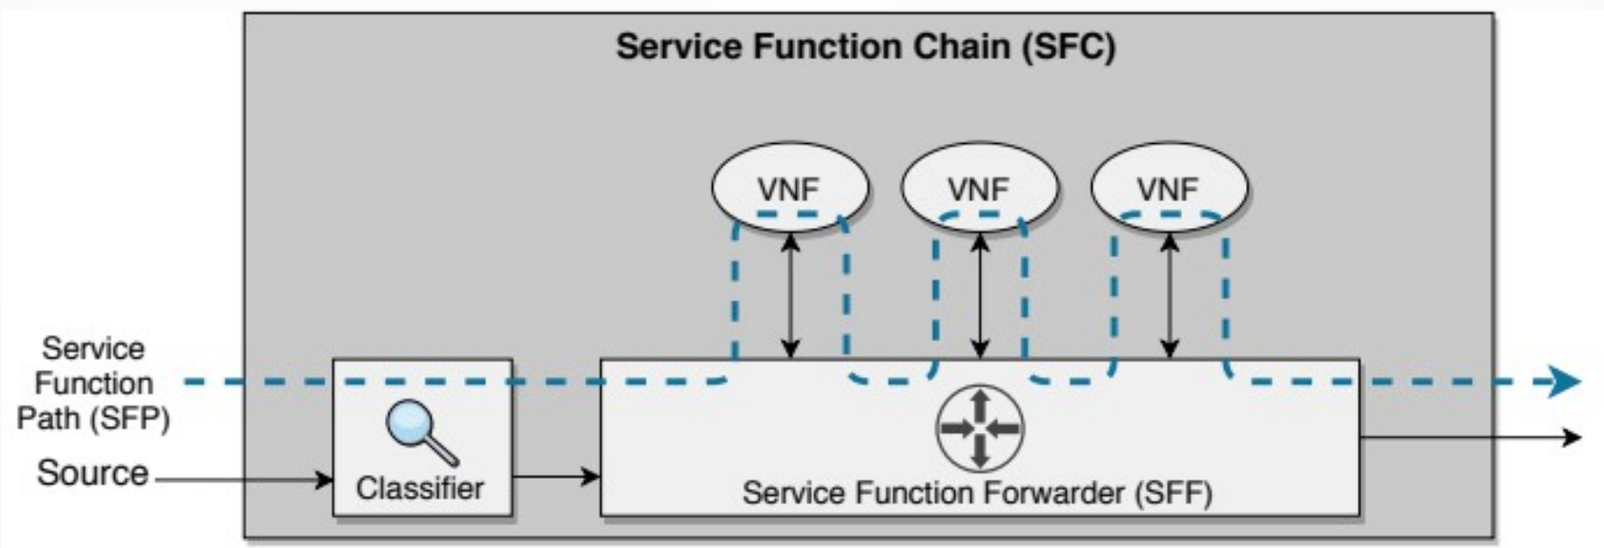

Different Types of Services: Different Requirements

- Stateless, Stateful per Flow, Stateful among Flows

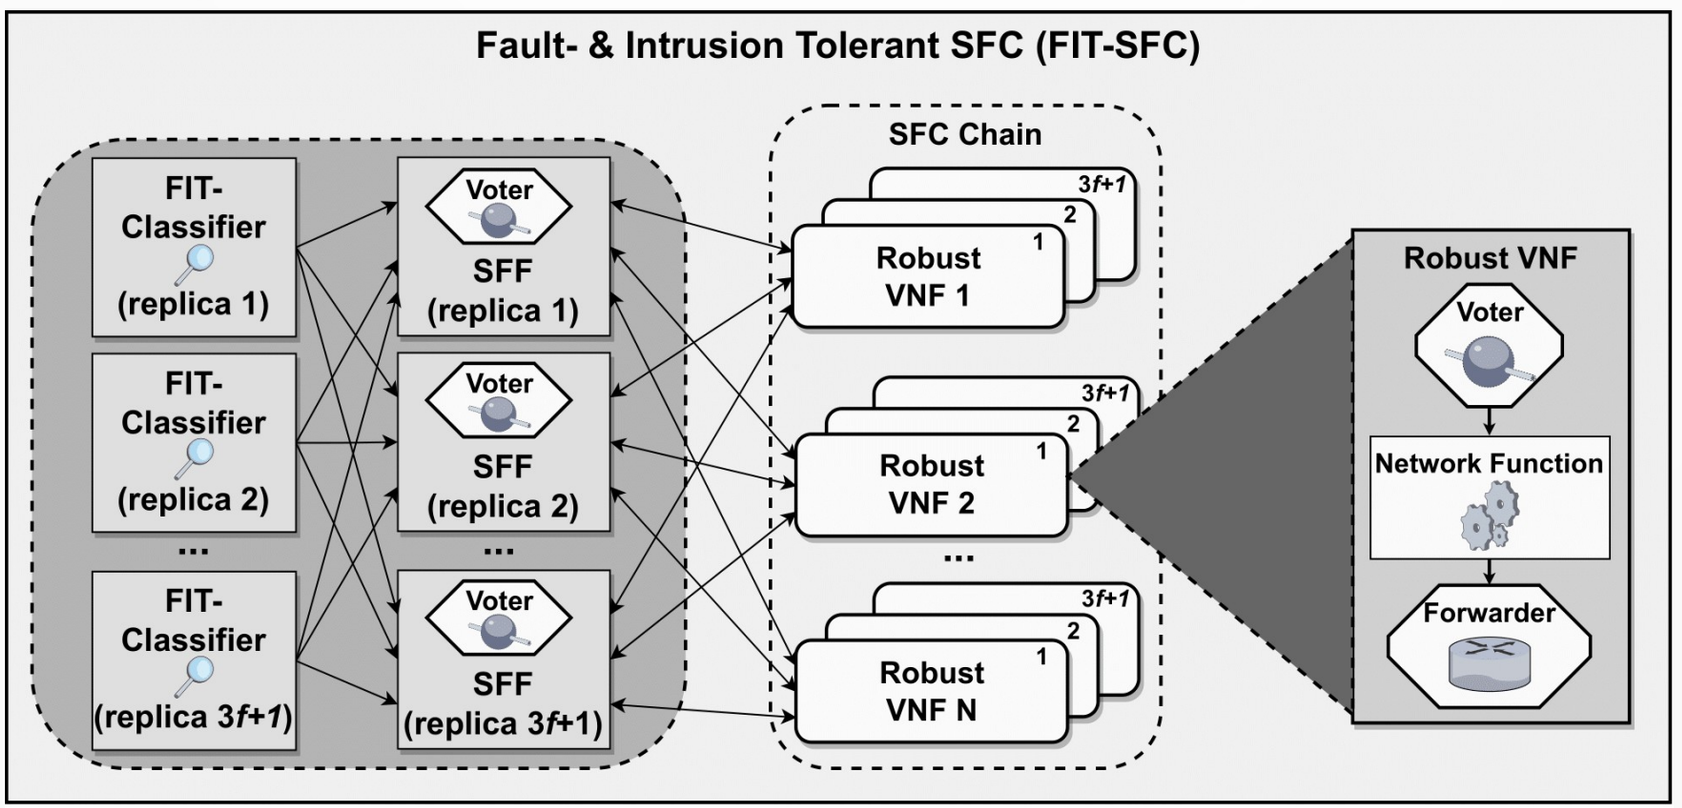

HyperPaxos: Hierarchical Consensus

- A scalable version of Paxos
- Acceptors form a vCube virtual topology
- vCube organizes processes of a distributed system
  - Scalable by definition: when all processes are correct → vCube is a hypercube

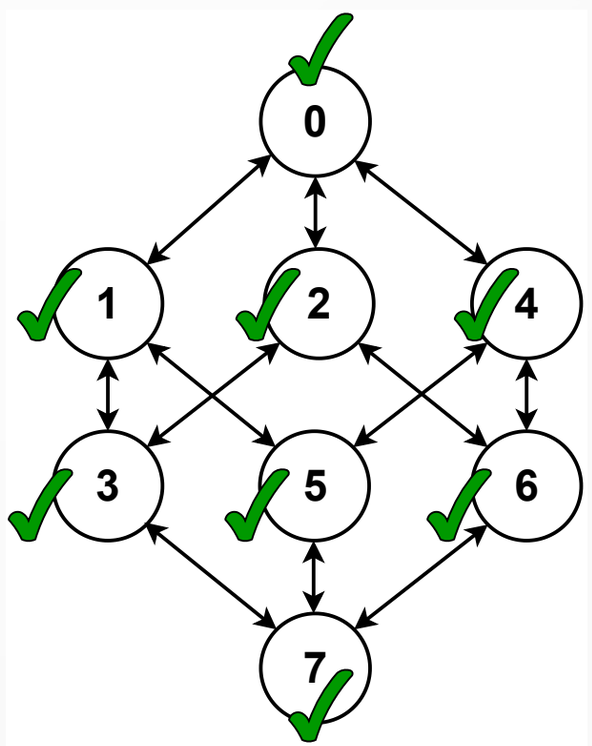

vCube: Scalable by Definition

- As processes crash: the topology autonomously reorganizes itself, maintaing several logarithmic properties

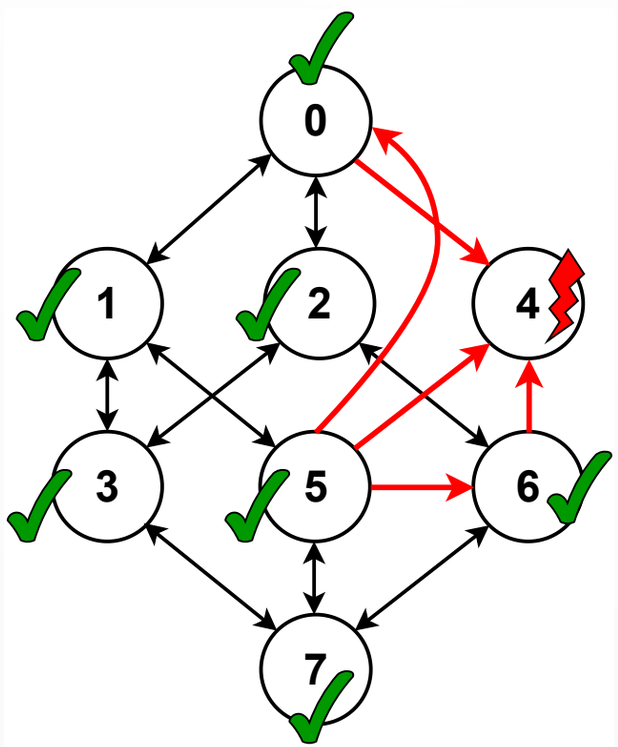

Alternatives to vCube: Fully Connected, Ring, Gossip

- Either too many messages or it takes too much time to communicate or probabilistic

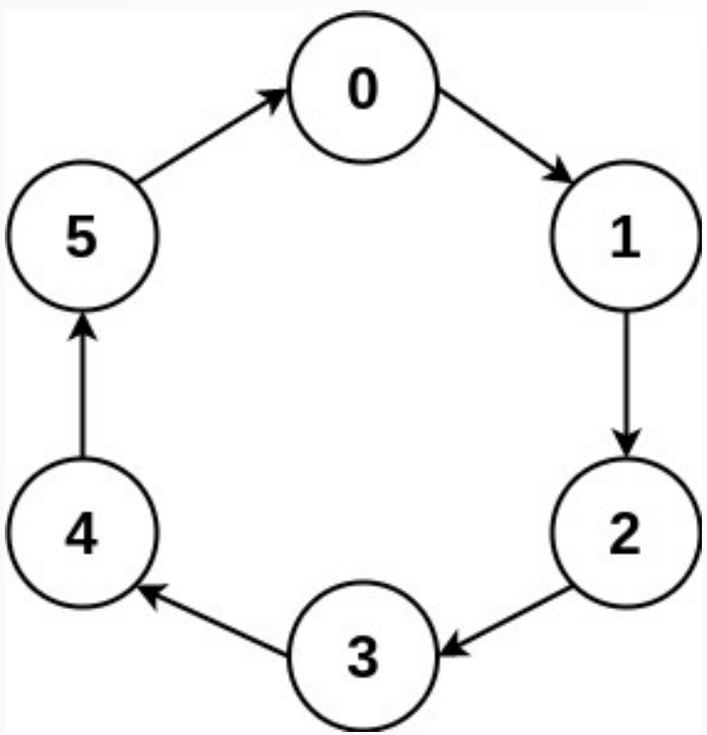

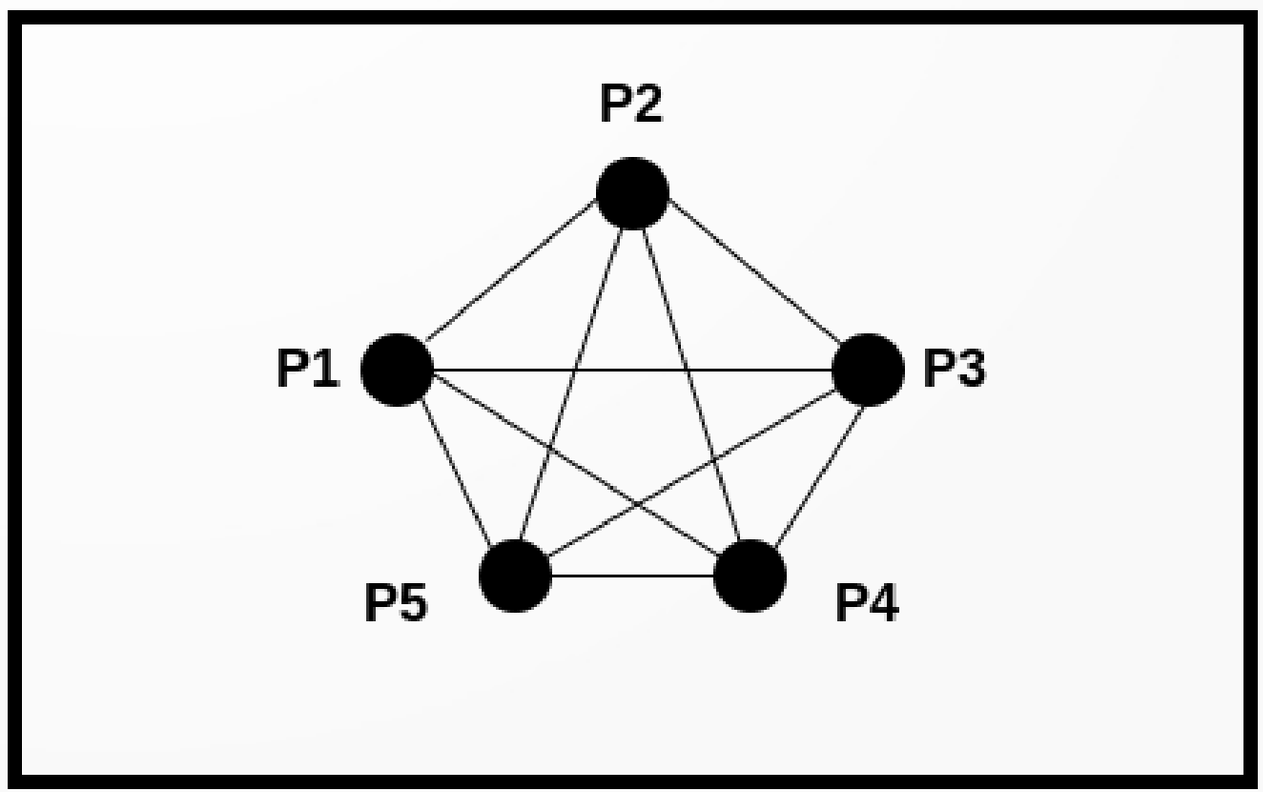

Conclusions

- The synergy of NFV & COIN: the network can support novel services in a flexible way
- FIT: Fault- and Intrusion-Tolerant Service Function Chaining
- HyperPaxos: Hierarchical, Scalable Consensus

Thank you!

elias@inf.ufpr.br
www.inf.ufpr.br/elias