future internet

MDPI

Article

Network Function Virtualization and Service Function Chaining Frameworks: A Comprehensive Review of Requirements, Objectives, Implementations, and Open Research Challenges

Haruna Umar Adoga * and Dimitrios P. Pezaros

School of Computing Science, University of Glasgow, Glasgow G12 8QQ, UK; dimitrios.pezaros@glasgow.ac.uk
* Correspondence: h.adoga.1@research.gla.ac.uk

Abstract: Network slicing has become a fundamental property for next-generation networks, especially because an inherent part of 5G standardisation is the ability for service providers to migrate some or all of their network services to a virtual network infrastructure, thereby reducing both capital and operational costs. With network function virtualisation (NFV), network functions (NFs) such as firewalls, traffic load balancers, content filters, and intrusion detection systems (IDS) are either instantiated on virtual machines (VMs) or lightweight containers, often chained together to create a service function chain (SFC). In this work, we review the state-of-the-art NFV and SFC implementation frameworks and present a taxonomy of the current proposals. Our taxonomy comprises three major categories based on the primary objectives of each of the surveyed frameworks: (1) resource allocation and service orchestration, (2) performance tuning, and (3) resilience and fault recovery. We also identify some key open research challenges that require further exploration by the research community to achieve scalable, resilient, and high-performance NFV/SFC deployments in next-generation networks.

Keywords: network function virtualization; service function chaining; software defined networking; next-generation networks; SFC frameworks

Citation: Adoga, H.U.; Pezaros, D.P Network Function Virtualization and Service Function Chaining Frameworks: A Comprehensive Review of Requirements, Objectives, Implementations, and Open Research Challenges. *Future Internet* **2022**, *14*, 59. https://doi.org/10.3390/fi14020059

Academic Editor: Eirini Eleni Tsiropoulou

Received: 21 January 2022
Accepted: 3 February 2022
Published: 15 February 2022

1. Introduction

With the advent of software-defined networking (SDN) and network function virtualization (NFV), middlebox functionality is increasingly being virtualized and provided in software, which can reduce power consumption, resource usage, and operational costs for service providers [1,2]. This paradigm is a departure from the use of hardware middleboxes, which are often proprietary, and thus not easily extendable by service providers. By abstracting network functionalities and implementing them in software, network operators can create network functions that suit their service level agreement (SLA) and service models [3].

NFV helps network operators by providing the right environment for the rapid deployment and scaling of virtual network functions (vNFs) [4], which are chained together in what is known as the service function chaining (SFC) architecture. Unlike the typical routing technique employed by middleboxes, where packets are simply forwarded directly from source to destination, SFC routes packets through a chain of network functions before reaching the destination (depending on the type of service and policy in use). Thus, one of the major goals of SFC is the flexibility it offers by allowing traffic to traverse diverse network functions along the service chain.

Heterogeneous network services deployed using SFC and virtual functions are installed on multiple virtual machines, sometimes combined with containers [5,6], which are chained together to provide services to users. SFC is generally considered as one of the important use cases of NFV and SDN architectures [7], which is also made possible

using a centralised network controller that has a *global* logical view of the entire network infrastructure, and handles tasks such as the creation of service chains and orchestration of traffic between vNFs [8]. In terms of the location of network functions in a service chain, the virtualized infrastructure can span multiple datacentres, which calls for inter-data-centre networking or within the same data centre, which results in an intra-data-centre network (see Section 2.2).

We generally expect the SFCs to be adaptive and dynamic; service chains should have the feature of readjusting to the unpredictable nature of service requests, and to offer better quality of service (QoS) for the end user [9,10]. If we consider network security as an example, such a service may consist of network functions such as firewalls, deep packet inspection (DPI), intrusion prevention or detection modules installed as software, which are chained together [11].

Several implementation frameworks have been proposed in the literature, some of which have been adopted and deployed at a commercial scale, while others are in the experimental state of development. In this paper, we put together these frameworks while highlighting their contributions and how they solve research problems associated with realising NFV/SFC in service provider network environments.

Experts from both academia and industry are putting a lot of effort in the research and development of SFC. While NFV/SFC has developed over the years, developments in this domain are so substantial that it makes a survey now quite different and valuable from a (hypothetical) survey put together a few years ago. Unlike previous surveys in this area, our work (layout in Figure 1) also shows how SDN, NFV, and SFC are combined in the provisioning, operations, and management of next-generation networks. We further explain how our work differs from previous efforts by presenting related surveys in Section 2; thus, our contributions are threefold:

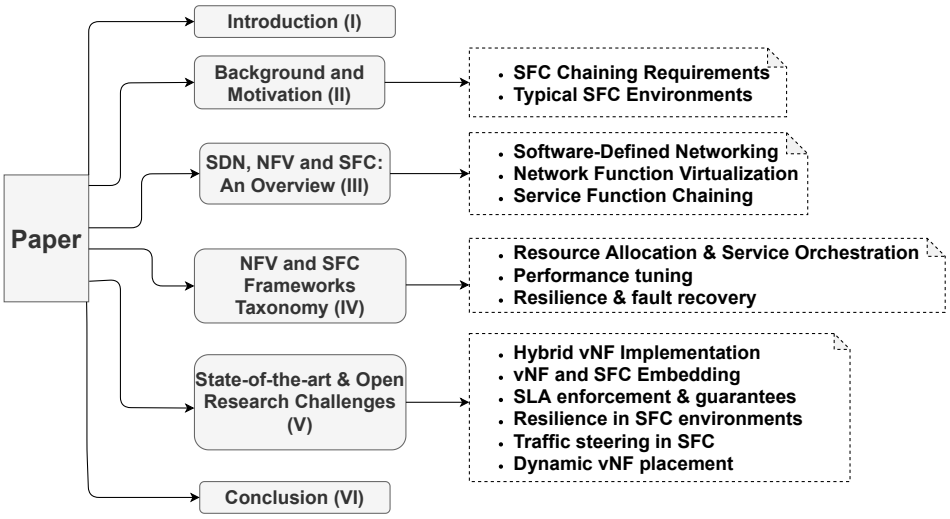

Figure 1. High-level paper structure.

- We present the vNF chaining requirements (Section 2.1) and how they relate to SFC implementations in service provider networks, thus some use cases for the SFC paradigm (Section 2.2) are presented. An overview of the current state-of-the-art SDN, NFV, and SFC (Section 3) is also presented as the foundation for the frameworks reviewed in this paper;
- We present a comprehensive survey of the state-of-the-art NFV frameworks for building and implementing vNFs, particularly frameworks that have been proposed for use in SFC environments. We also present a taxonomy of SFC implementation frameworks, focusing on each SFC-related challenge, each one addressing the approaches used (Section 4);

- Finally, we identify and discuss the main open research challenges associated with NFV and SFC environments in the next generation networks (Section 5).

The remainder of this paper is structured as follows: Section 2 presents the background and motivation of this survey, considering what has been done in the literature and highlighting the uniqueness of our approach and reflections on the state of the art approaches. Section 3 briefly discusses the underlying technological advancements that make NFV possible. Section 4 contains a taxonomy of the implementation frameworks considered in this work, classifying the frameworks into three major categories. The problems addressed by the frameworks are presented in Table 3 (resource allocation and service orchestration), Table 4 (performance tunning) and Table 5 (resilience and fault recovery) which represent various aspects of the SFC implementation challenges. We present the state-of-the-art SFC and highlight the open research challenges of NFV/SFC environments in Section 5. Finally, Section 6 concludes the paper.

Research Methodology

Our work covers research performed in recent years in the area of NFV/SFC implementation frameworks. The methodology employed involved a critical review of the relevant contributions in report papers, journals, conference papers, and articles. We systematically queried the most relevant scholarly databases (IEEE, ACM Digital library, Web of Science, Scopus, etc.) to get all the required research papers that are relevant to our work. We started this process with 367 papers, of which 118 papers met our final selection criteria.

We made use of some basic exclusion criteria (Figure 2) such as the exclusion of papers that were written in languages other than English, papers with no full text, research papers in duplicates, the year of publication, and the specific area(s) covered such as focusing on research papers in areas related to NFV, SFC, and SDN in service provider network environments.

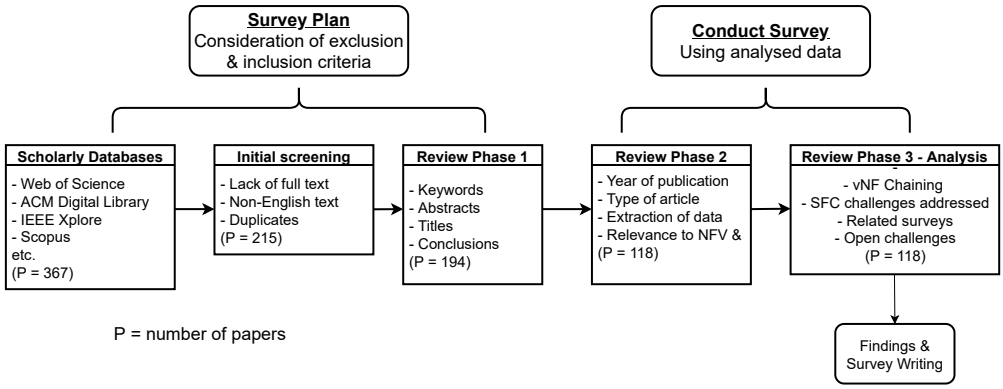

Figure 2. Paper analysis process.

For the review of papers presented in our proposed taxonomy, we included research papers that present NFV implementation frameworks with some vNF chaining components, representative frameworks that address specific problems in the domain, e.g., vNF orchestration, traffic steering, SFC resilience, modular SFC deployment, etc. A total of 118 papers were finally selected, which include some notable related surveys, and experimental and theoretical works in NFV/SFC. The papers we have cited in our taxonomy cover fundamental equivalence classes of the state-of-the-art in NFV/SFC implementation frameworks.

In addition to the more recent research papers presented in our work, we would like to note that some of the papers we considered present both the pioneering efforts for certain equivalence classes of NFV/SFC frameworks and still represent the state-of-the-art; thus, the most recent works (if any) are incremental to what we have presented. As a concrete example that justifies some of our choices (based on current literature), frameworks such as ClickOS [12], OpenNetVM [13], OpenBox [14], Slick [15], SNF [16],

Metron [17], NetBricks [18], and GNF [19] constitute some of the notable pioneering efforts in this domain, which are still relevant in designing NFV/SFC frameworks for next generation networks.

Each of these frameworks represent important aspects such as the use of lightweight packet processing elements, the use of packet acceleration frameworks, the decoupling of network infrastructure into control and data-planes, synthesizing packet processing elements for improved performance, resource allocation, and deploying containerized network functions at the network edge to support the edge computing paradigm. In terms of technological advancements and existing works in the literature, we have covered the core components that are still used in designing and implementing NFV/SFC frameworks to the best of our knowledge.

2. Background and Motivation

There are several existing studies that focus on the general areas of SDN, NFV, and SFC, as summarised in Table 1. Each of the related surveys presented in our work explores different aspects such as NFV implementation problems, traffic steering in SFC, NFV deployment acceleration, optimal NFV and SFC concepts, design, and taxonomy of NFV platforms, the placement of network functions, and resource allocation. This section summarises the key contributions made by existing related surveys and provides the foundation for our research.

Table 1. Existing NFV/SFC Surveys.

Survey Reference	Survey Focus	Open Challenges Presented	Year
Zoure et al. [20]	Network services anomalies in NFV	✔	2022
Zhang et al. [21]	NFV platforms, design and taxonomy	✔	2021
Hamdan et al. [22]	Load balancing techniques in SDN	✔	2021
Fei et al. [23]	NFV deployment acceleration	✔	2020
Kaur et al. [24]	vNF placement, availability, and load balancing	✔	2020
Hantouti et al. [25]	Service function, 5G and next-generation networks	✔	2020
Bonfim et al. [26]	NFV architectures; NFV/SDN design taxonomy	✔	2019
Hantouti et al. [27]	Traffic steering in SFC; SDN-based chaining evaluation	✔	2019
Laghrissi et al. [4]	Service placement survey; vNF placement; existing virtual resource placement solutions	Lessons learned	2019
Mirjalily et al. [28]	Optimal NFV and SFC concepts	✔	2018
Medhat et al. [29]	Next Generation Network SFC; SDN approaches in SFC; SFC implementations	Implementation limitations	2017
Bera et al. [30]	SDN for IoT	✔	2017
Veeraraghavan et al. [31]	NFV survey; selected implementation problems and solutions	✘	2017
Bhamare et al. [32]	Service function chaining; NFV mobility	✔	2016
Herrera et al. [1]	NFV resource allocation	✔	2016
Xie et al. [33]	Resource allocation in SFC; Existing NFV SFC RA solutions	Lessons learned	2016
Li et al. [7]	Network function placement; Selected framework comparison	✔	2016
Yang et al. [34]	SDN/NFV for mobile and wireless networks	✔	2015

As one of the most recent survey works in the NFV/SFC space, Fei et al. [23] focused their work on the proposals that consider the acceleration of NFV deployments. A taxonomy of the surveyed approaches is presented and discussed, which mainly involves the hardware and software acceleration of NFV deployments. We take a different approach by considering the state-of-the-art NFV implementation frameworks, especially in scenarios where network functions are also chained, as in SFCs.

In their work, Zhang et al. [21] presented NFV platform design choices. They presented three main open issues in NFV: (1) the use of artificial intelligence in NFV, (2) network slicing, that is, the management of network slices, the communication between slices and placement of network slices, and (3) the integration of NFV with IoT. The taxonomy of NFV platforms presented by Zhang et al. consists of prototyping, testing, deployment, management, execution, and integrated NFV platforms. We take a different approach in our work by first presenting the requirements that must be met for the chaining of vNFs and some useful use-cases. This is relevant in understanding the key components involved in creating service function chains in service provider network environments. Our work

focuses mainly on NFV frameworks that employ the chaining of network functions for effective service delivery. In addition to the taxonomy of the frameworks presented, we also discuss key open research challenges in these environments.

In the survey presented by Hantouti et al. [27], they discussed the traffic steering approaches used in SFC solutions with SDN. The current traffic steering approaches are classified into three methods: header-based, tag-based, and programmable switch-based methods. They concluded by identifying QoS, scalability, security, and management as some of the challenges in SFC traffic steering. Although their survey is comprehensive in terms of SDN-based traffic steering in SFC, we focus our work not only on traffic steering in SFC, but also on the implementation frameworks used to achieve a scalable, resilient, and high-performance SFC.

Bonfim et al. [26] presented a review of integrated NFV/SDN architectures, where they considered implementation frameworks that combined NFV and SDN. They identified vNF scheduling and placement, improving network programmability, and the possibility of deploying multiple SDN controllers to achieve scalability, security, and standardisation of SDN/NFV solutions as some of the open challenges. We extend their contributions by considering proposals that implement NFV/SFC and the challenges that require further research.

As one of the early attempts in this domain, the survey by Yang et al. [34] explores the challenges faced in mobile and wireless network (MWN) environments. Their work carefully describes how software-defined wireless networks (SDWNs) and wireless network virtualization (WNV) can be used in addressing the challenges of MWN networks.

A survey on SFC was presented by Herrera et al. [1] and Xie et al. [33], where the authors focused on resource allocation approaches in the literature. Medhat et al. [29] presented open challenges in service function chaining for next-generation networks. A taxonomy of prior work was presented, where they classified it into data -and control-plane SFC solutions. Our survey considers a broader scope by creating a taxonomy that is beyond control and data-plane solutions, as well as presenting state-of-the-art technologies and open research challenges.

Bhamare et al. [32] presented a detailed survey of SFC, where the authors identified optimal resource allocation, dynamic service mapping, and policy enforcement, as some of the challenges in these environments. Li et al. [7] and Laghrissi et al. [4] focused on the placement of resources in SFC environments, where the former presented a survey of network function placement in SFC, and the latter focused on the placement of virtual resources. In their survey, Hamdan et al. [22] explored the traffic load balancing approaches used in SDN network environments. Our work presents the relationship between SDN, NFV, and SFC, as well as a novel taxonomy of implementation frameworks.

Mirjalily et al. [28] presented a survey of SFC and NFV implementation efforts. Some of the future research directions presented by [28] include SLA and QoS approaches, online chaining of service functions, availability and resilience of chains, security, and energy efficiency. Our survey provides a more technical look into the implementation frameworks by first classifying the state-of-the-art frameworks and the proposed solutions.

The comprehensive survey presented by Bera et al. [30] focus on SDN technologies employed in network environments such as the data centre, edge, access, and core networks. They present their findings in relation to IoT use cases. They identified open research challenges such as platform independence, policy enforcement, mobility management, and the fully practical implementations of SDN-based solutions in IoT environments.

In their comprehensive survey, Kaur et al. [24] classified SDN/NFV approaches into availability, placement, and load-balancing solutions. They identified the ordering of SFCs, resiliency, security, topology configuration, and service placement, as some of the challenges related to current SFC implementations. We extend this work by (1) creating a taxonomy that captures state-of-the-art frameworks; (2) presenting frameworks that also capture the chaining of vNFs, which is key to the design of next-generation networks;

and (3) presenting more open challenges that are key to achieving scalability of network functions, resilience, and high performance in NFV/SFC network environments.

The survey carried out by Hantouti et al. [25] on SFC challenges covered frameworks that proposed to solve challenges such as path selection, orchestration, security, SFC path composition, QoS, and traffic steering. They concluded by acknowledging that more work needs to be done in developing related technologies (SDN and NFV). We also consider frameworks that focus on areas such as resilience, fault recovery, performance tuning, and resource allocation to extend their contributions.

To the best of our knowledge, we have summarised the key contributions of notable related surveys and how our work adds to the NFV/SFC domain. Unlike most existing surveys, our research also presents the detailed requirements for chaining virtual network functions in service provider environments, including some important use cases. The taxonomy we created in our work presents the design choices and technologies used for the implementation of different equivalence classes of NFV/SFC frameworks. This focuses on the existing problems in the NFV/SFC domain, that is, we categorise the surveyed frameworks based on their proposed solutions and technological design choices.

Most of the notable related surveys we have presented only discuss the technologies and open challenges in the NFV or SFC domain. Some exceptions to this are the surveys by Bonfim et al., which present a taxonomy of NFV/SDN architectures by categorising them into NFV-side and SDN-side designs, while Zhang et al. presented a taxonomy based on the life cycle of NFV platforms. Kaur et al. presented a taxonomy based on the optimisation approaches used in SFC, such as availability, placement, and load balancing. In contrast to previous surveys, we categorise the surveyed frameworks based on their proposed solutions and technological design choices.

We carefully describe what each framework has been designed to achieve, what technology and approach have been used, and what performance (or other quality) benchmarks have been performed. We have also highlighted the different methods and technologies used for implementing each framework. In addition to presenting an extensive list of open research challenges in NFV/SFC, we also discuss some notable early attempts in the literature aimed at addressing the research challenges identified (Section 5), providing the reader with knowledge of some existing efforts in this direction.

2.1. SFC Chaining Requirements

When it comes to the chaining of network functions in an SFC environment, the IETF SFC draft [35] provides an architectural framework that captures all the components that are required for SFC implementation in service provider networks.

There are some useful assumptions that need to be considered when creating a chain of network functions: (1) different network functions present their own configuration and description challenges, thus, creating a generalised description for all service functions is not trivial; (2) the implementation environment of the network function affects the list of functions that can exist in a particular domain; (3) the logic employed for the chaining of service functions is not fixed, that is, it is peculiar to any given administrative domain and the requirements of service(s) to be delivered to end users; and (4) the invocation of any service chaining criteria depends on the administrative domain in which the service functions are deployed [35].

In terms of chaining requirements, although there are domain-specific requirements that need to be in place when deploying SFCs (based on the network administrative domain), there are also general requirements for the components that are found in most SFC deployments [36]. Irrespective of the network environment(s), *global* components must be in place. These components are the service classifier, which is placed at the entry point of the network (to classify ingress flows). Flow classification helps with the decision-making process of the orchestrator in terms of traffic steering across the service chain.

The service function forwarder (SFF) is another component that forwards received traffic to the right service function (SF), and can be embedded on a physical network

component or deployed as a virtual component along the service function path (SFP), which is based on the classification of ingress traffic performed by the SC [37]. The SFF is also responsible for handling any return traffic that needs to be forwarded back to a specific service function or service classifier in the service chain [35].

Another component is the SFC proxy, which is often optional in SFC implementations and used in scenarios where other components (SF and SFF) are unable to communicate in the chain [25]. SFC deployments can be fully deployed without the use of any form of proxy component, that is, when SFC-unaware service functions are not deployed in the network infrastructure. In situations where SFC-unaware service functions are part of the service chain, an SFC proxy is used to add or remove encapsulation information; thus, these are considered as *logical* components of the SFC architecture [35]. Figure 3 depicts some of the core components of the IETF SFC architecture.

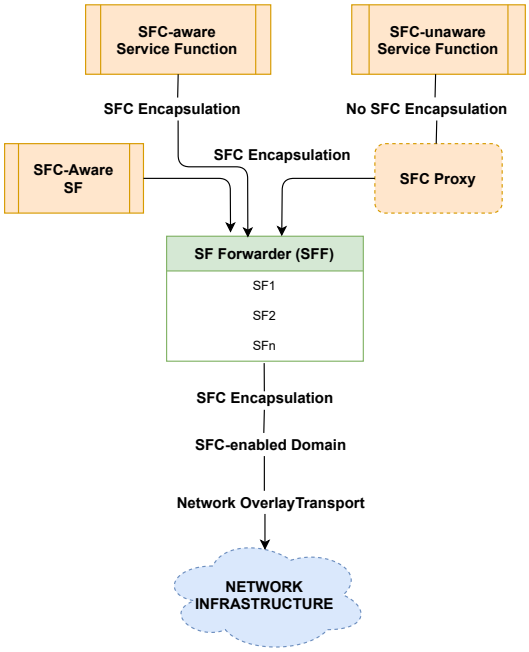

Figure 3. IETF SFC Architecture Components.

2.2. Typical SFC Environments

As more network operators continue to adopt network slicing, which serves as the enabler for next-generation networks, the chaining of virtual network functions for efficient service delivery has become commonplace. Some common use cases can be found in today's service provider networks. Some available common environments are the Gi-LAN network used by mobile network operators, residential/consumer services, and inter/intra-data-centre networks. Mobile network operators deploy functions such as traffic optimizers, firewalls, carrier-grade network address translation (NAT), load balancers, and DPI, at the core of the network, which is designed for subscribers that access Internet-based services [38]. Here, we briefly explore these environments.

2.2.1. The Gi-LAN Mobile Core Network

A typical environment in which service function chaining is deployed is the Gi-LAN network, which is a component of mobile networks used by operators to provide fine-grained user-specific services such as traffic optimisation, DPI, and firewalls [39]. Gi-LAN implementation by mobile network operators is an emerging use case for SFC architecture [40]. These services are often chained together and are provided by multiple vendors, and the traffic is steered to the right service functions, which is aimed at meeting service level agreements and policy enforcement [41].

The ability to add or remove a service becomes easier along the processing pipeline, as per-user services can be created for mobile data monetisation [42]. In terms of implementation requirements, implementing a service chain that contains a network function such as NAT, for example, requires that the function is placed on the edge sites, which is closer to the users requesting the service [43]. Another useful requirement when creating a service chain in a mobile core environment is the consideration of the SLA agreement between multiple vendors. For packet classification requirements, the classifier should be located at the packet gateway. Figure 4 depicts a high-level description of the SFC in mobile networks, with service functions deployed in the Gi-LAN segment of the network.

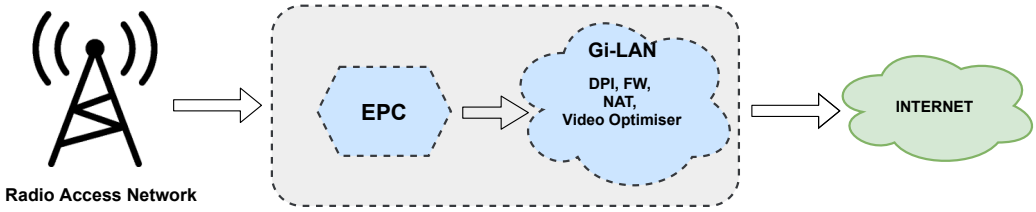

Figure 4. Gi-LAN in mobile networks.

2.2.2. Residential and Consumer Services

Because NFV/SDN allows for the provisioning of highly specialised solutions to meet customers' quality of service requirements, using the concept of service function chaining, service providers can steer residential traffic such as parental control and VoIP-related services. The idea of follow-the-user service deployment in residential environments is an important use case that is achieved using SFC implementations.

Network operators make use of vCPEs to easily create a chain of services that meet user requirements [44]. Users in these environments are more likely to make use of web-based applications that use HTTP as the de facto protocol [43]. One of the requirements for such deployments is to create a service chain that prioritises security [45]. A typical example is a service chain which follows the order: firewall > IDS > proxy. Security of users is a priority and an important requirement, especially in this scenario.

2.2.3. Inter-and Intra-Data-Centre Networks

SFC in the inter/intra DC environment allows for the chaining of virtualized enterprise network applications (in the case of intra-data-center) and chaining across multiple locations, or inter-cloud, in the case of inter-data-centre networking. The ability of service functions to be instantiated across multiple datacenters (inter-data-centre networking) is a key requirement for live VM migration. The SFC architecture to be implemented should be designed to dynamically migrate service functions from one VM/container to another without disrupting user service requests [46].

Deploying and managing service function chains in an inter-data-center setting incurs inter-data-centre bandwidth cost, deployment cost, intra-data-centre cost, and vNF costs [47]. In these environments, NFs are also used for policy-based routing of cloud services and enterprise applications [38], which means that service providers can up-sell their services easily using SFC at the enterprise. This can be achieved by making network services user-programmable.

3. SDN, NFV, and SFC: An Overview

In service provider networks, the process of creating, deleting, modifying, and steering traffic in SFCs is carried out efficiently by using SDN and NFV technologies [48]. These technologies are the key networking paradigms that are at the core of the frameworks surveyed in our study. In this section, we describe these technologies as they relate to NFV/SFC implementation frameworks in service provider network environments, thus showing their interrelation in the operations of next-generation networks. Even though the chaining of hardware middle-boxes is possible, the use of NFV makes it much easier

and cheaper [49]. Thus, SDN is employed for orchestrating virtual network functions by providing a centralised logical control and the creation of service chains.

3.1. Software-Defined Networking

Software-defined networking decouples the control plane from the data plane in the networking devices. Traditional non-SDN networks often have control and data planes integrated on a single device, which brings about challenges such as management complexity and scalability issues. Implementing centralised network control using SDN controllers results in easier service deployment and management [49]. This helps service providers to easily steer traffic between NFs by scaling across multiple physical machines.

The functional separation of the network infrastructure into control and data planes, as shown in Figure 5, is the core concept behind SDN. The application layer consists of various network applications, providing network services that use the Northbound Interfaces for sending requests to the control plane (centralised logical control). A global view of the network infrastructure is maintained by an SDN controller such as OpenDaylight [50], POX [51], RYU https://ryu-sdn.org/ (accessed on 12 December 2021), or a custom-built controller can be used to manage network functions, which handles requests coming from the network applications, and sends instructions to the data plane of the network for packet processing [52].

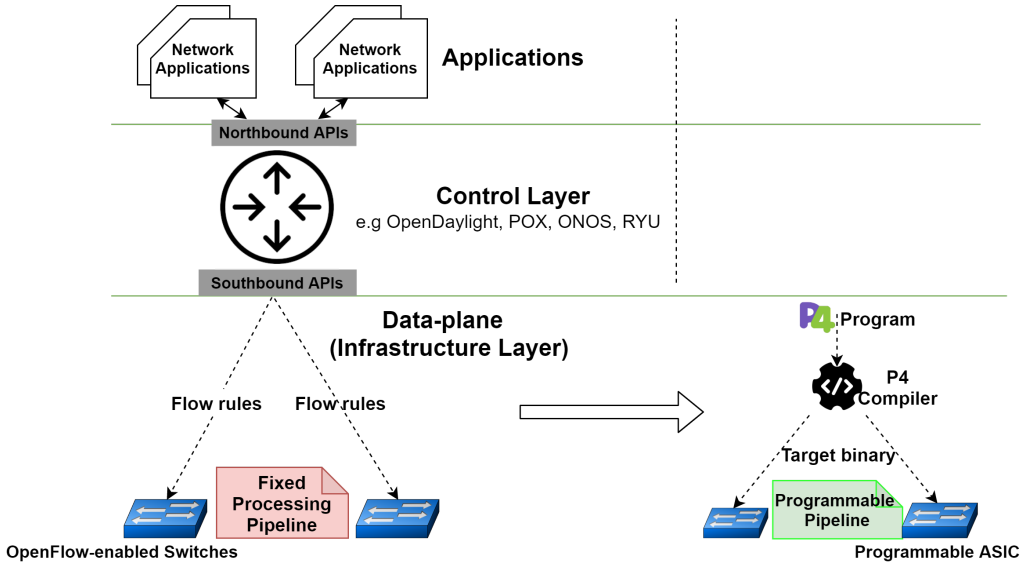

Figure 5. OpenFlow SDN and P4 data-plane.

Using the southbound application programming interfaces (APIs) and the OpenFlow protocol, rules are sent down to devices in the data plane of the network, which is responsible for packet processing and forwarding. Although traditional SDN networks use the OpenFlow protocol to communicate with the network data plane by inserting flow rules on devices, the network has become more programmable over the years. Programmability allows network operators to define the processing pipeline and how packets are processed using high-level languages such as P4 [53].

Figure 5 depicts a high-level comparison of the operation of OpenFlow SDN and programmable data planes, where P4 programs are written and compiled and then deployed on programmable switches. This allows for the creation of a highly programmable pipeline, as opposed to a fixed (less flexible) processing pipeline found in OpenFlow SDN. SDN and NFV serve as the building blocks for reaching the goal of deploying a chain of virtual network functions in the service provider network environment.

3.2. Network Function Virtualization

The use of proprietary network hardware is expensive for service providers in terms of procurement, security, configuration, scalability, and maintenance costs. The European Telecommunications Standards Institute (ETSI) [54] introduced a high-level NFV architectural framework, envisaging the deployment of network functions as software, running on the network function virtual infrastructure (NFVI), which could be a general-purpose server. This proposal was introduced to take advantage of hardware virtualisation [55,56]. The deployment of network services has been greatly simplified by NFV, because the cost of acquiring new hardware middle-boxes is reduced, and several middle boxes can be virtualized and deployed on single or multiple general-purpose servers.

The ETSI architectural framework for NFV depicted in Figure 6 shows all the important components that are necessary for deploying NFV. The operations support system (OSS) and business support system (BSS) directly interact with the vNFs. The vNF component is the network functionality, for example, a traffic load balancer, a WAN optimizer, and a firewall, etc. (Table 2). The hardware infrastructure consists of a virtual infrastructure with virtual computing, storage, and network components. This infrastructure is managed by the virtual infrastructure manager (VIM), which is responsible for resource allocation and embedding of virtual network functions on the virtual infrastructure. NFV Orchestrator (NFVO), which is an integral part of the ETSI NFV framework [54], is responsible for service orchestration and management [29]. One of the functions of the orchestration layer is the mapping of virtual network functions in a service chain to available physical resources. The management and network orchestration (MANO) component is responsible for the orchestration of vNFs and the chaining of services in a scenario where service function chaining is used. As shown in the SFC scenario in Figure 7, general-purpose hardware could be a typical high-performance commercial off-the-shelf (COTS) hardware [54].

Table 2. Commonly used vNFs in NFV.

vNF	Functionality
Application gateway	Layer 7 traffic management based on application profile
Layer 2 forwarder	Packet forwarding based on layer 2 information
Protocol Analyzer	Packet classification, based on protocol in use
Flow tracker	Storing, displaying and forwarding ingress flows
Layer 3 Switch/Router	Traffic routing and switching using IP addresses
Application Firewall	Layer 3 and layer 7 packet filtering
Bridge	Bridging between two networks or host devices
Carrier Grade NAT	IP address translation for WAN connectivity
IDS/IPS	Stateful or stateless intrusion detection and prevention
Protocol converter	Protocol translation between IPv4 to IPv6
Encryption gateway	Packet encapsulation and packet encryption/decryption
ACL	User and application level access control
Protocol Accelerators	Performance improvements by ISPs
VLAN manager	VLAN encapsulation and decapsulation

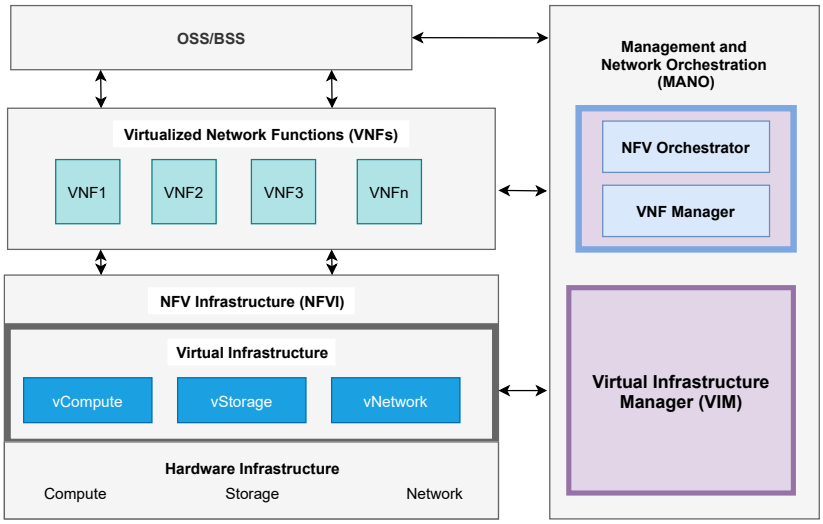

Figure 6. ETSI NFV architectural framework.

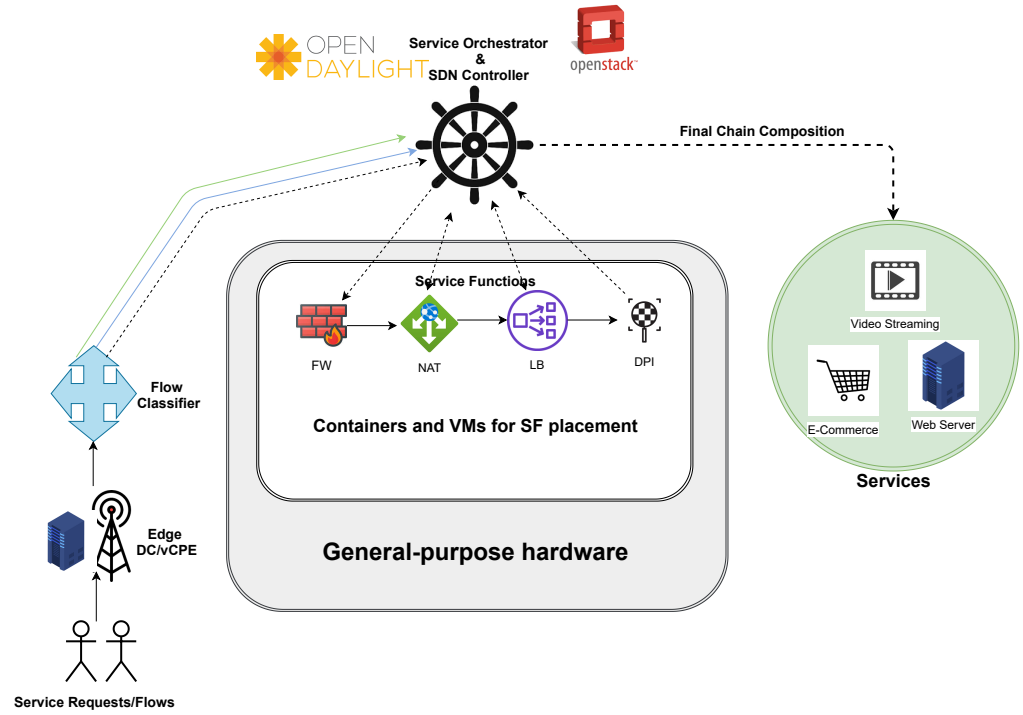

Figure 7. Service function chaining scenario.

3.3. Service Function Chaining

The vNFs implemented in NFV constitute the NF forwarding graph, which consists of network functions connected via logical network links to achieve the goal of packet processing by the vNFs. SFC is made up of NFs connected in a chain (based on service requirements and specifications) to deliver end-to-end services to end users [54,57]. A typical service chain consists of NFs such as a NAT function, a firewall, and a traffic load balancer. As briefly introduced in Section 2.2, in order for an SFC deployment to be complete, components such as the service classifier, SFF, Service Function Path, SFC proxy, and service function need to be in place.

The IETF SFC architecture presented by Halpern et al. [37] shows that the SFCs are either bidirectional or unidirectional, where packet processing is performed through an ordered list of service functions in a unidirectional scenario [37]. A bidirectional SFC scenario requires packet processing elements (SFs) to be placed in both directions of the service chain. SFCs are deployed as network service graphs with SFs placed carefully

at different parts of the service chain. The ability to add and remove SFs dynamically along the service path is essential for the design of any SFC framework. NFV and SDN are integrated to achieve instantiation, management, and orchestration of service chains [58].

Intelligent service orchestration is important when handling various service functions, and this can be achieved when the NFV is properly integrated with SDN [59]. Figure 7 depicts a typical SFC scenario with service requests generated by users based on the application requirements. A classification of the user traffic is carried out by the flow classifier, which helps in deciding what network function(s) need to be traversed by the traffic before reaching its destination. In a scenario where the service requests traverse more than one NF, a service orchestrator is used to create a chain of NFs that forms the final processing pipeline toward the destination (requested service).

3.4. NFV/SFC and 5G Networks

The chaining of virtual network functions for effective end-to-end service delivery is a key enabler of Beyond 5G networks [60,61]. Since 5G-enabled networks are characterized by low latency, programmability, and the support for diverse use-cases of the future, technologies such as NFV can allow providers to deploy services that are suitable for radio access networks (RANs) and mobile core networks [60].

By implication, using NFV, SDN, and SFC, service providers can easily provide tailored solutions that meet customer demands, by carefully orchestrating user-generated traffic between an ordered list of network functions. As described in Section 2.2, the chaining of virtual network functions in SFC enables use cases such as the Gi-LAN mobile core network, residential and customer services, and inter and intra-datacentre networks. Other important use cases such as self-driving cars, e-healthcare [62], and mixed reality (MR) and 5G-enabled IoT [62,63] are also possible due to the flexibility offered by 5G network slicing.

Efforts such as the work by Morocho et al. [64] focus on showcasing how machine learning (ML) can be used to leverage the benefits provided by Beyond 5G networks. ML can be used with enhanced mobile broadband (eMBB) and support future Beyond 5G applications, that are envisaged to have high data rate requirements. Massive machine-type communications (mMTC) and ultra-reliable low-latency communications (URLLC) are also required to provide support for future use cases for Beyond 5G networks [65,66].

Abdelwahab et al. [67] explored how the 5G RAN can be enhanced using NFV, which could also lead to a reduction in the overall capital expenditure for telecommunications service providers (TSPs). As detailed in [67], some challenges that are related to 5G networks such as efficient scalability of vNFs between physical networks, vNF performance guarantees, and simultaneously supporting the deployment of hardware and virtualized network functions, can be overcome with the flexibility offered by NFV implementations.

4. NFV and SFC Frameworks Taxonomy

In this section, we present the implementation frameworks proposed for NFV/SFC deployments.

Some of these implementations are set out to solve specific problems in the SFC domain, such as resource allocation and service orchestration, performance tuning, resilience, and fault recovery. In Section 5, we present and discuss the open challenges that are related to the implementation frameworks discussed.

Figure 8 depicts a summary of the taxonomy of the frameworks presented in our work. Resource allocation and service orchestration frameworks deal with the efficient utilization of available resources, by employing techniques such as synthesizing packet processing graphs and offloading packet processing tasks onto smart NICs, while ensuring that traffic is steered to the right network functions in a service chain, and efficiently managing the life-cycle of network functions.

The frameworks presented under the performance-tuning category are concerned with improving the overall performance of SFCs by employing techniques such as modular SFC deployments, the use of lightweight packet processing elements, the use of acceleration

frameworks for packet processing, and deep learning techniques to improve the overall chain-wide performance.

The third category in our taxonomy, resilience, and fault recovery, consists of frameworks that handle the problem of fault tolerance in SFC. These frameworks employ techniques such as network function replication and piggybacking of NF state changes across service chains to achieve resilience in SFCs.

Each of the presented frameworks strive to achieve diverse objectives and are hence very often evaluated against different and incompatible baselines. We have therefore chosen to create a taxonomy of the surveyed frameworks to be able to compare, quantitatively and qualitatively, the different works in their own contexts. Rather than forcing a comparison of potentially disjoint performance characteristics between frameworks that could have led to superficial superiority claims, we have carefully described what each framework has been designed to achieve, what technology and approach has been used, and what performance (or other quality) benchmarks have been performed. Frameworks under each category are discussed next.

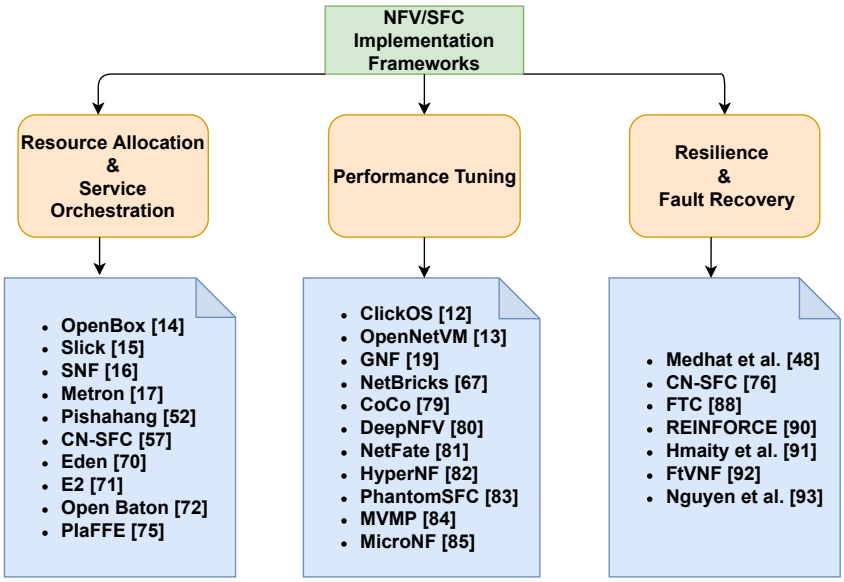

Figure 8. NFV/SFC frameworks taxonomy.

4.1. Resource Allocation and Service Orchestration Frameworks

The ability to efficiently allocate network, storage, and processing resources in virtualized network environments is key in service provider networks. This is even more so in NFV/SFC environments, where user traffic type and frequency can either be deterministic or nondeterministic, bringing the need for the underlying system to provide and allocate resources efficiently. The chaining of network functions to provide end-to-end services cannot be achieved without an efficient service orchestration scheme in place [68]. Efforts such as the work by Sun et al. [69] propose algorithms to handle SFC orchestration, resource utilization, and optimization. The algorithm proposed by [69] for example, is based on the breadth-first search (BFS) algorithm, which reduces overall chain-wide latency and bandwidth consumption. This section presents frameworks that have been designed to achieve the goal of efficiently allocating resources in NFV/SFC and the orchestration of network services. We categorically focus on frameworks that handle resource allocation and service orchestration as the core contributions in this subsection of our work by presenting the technologies used and the implementation approach employed by the authors. This has also been summarised in Table 3.

Table 3. Resource allocation and service orchestration frameworks.

Framework	Problem	Solution(s)	Processing	Testbed/PoC
OpenBox [14]	vNF provisioning and orchestration	Lightweight vNFs and decoupling control/data plane.	Click	Mininet
Slick [15]	Resource management and efficient traffic steering	vNF placement.	Click	Mininet/VM
SNF Framework [16]	Resource optimization	Synthesised processing graphs, stateful NF management and vNF performance optimization.	Click elements	Servers with MoonGen Packet gen
Metron [17]	Resource Allocation	Network resource management and server resource utilization.	Click-based	OpenFlow Switch and VM
Eden [70]	Service Provisioning and Orchestration	End-host NFs, packet offload to NIC and programmable data-plane vNF.	F# vNFs	End-hosts and programmable NICs
E2 Framework [71]	NF Scheduling and management	NF placement, effective resource allocation, vNF scaling.	commodity servers	Hardware switch and servers
Open Baton [72]	NFV/SFC management and orchestration	Network slicing, multi-site orchestration, vNF fault tolerance and resource allocation.	VMs	OpenStack
Pishahang [73]	Service Orchestration	Multi-domain orchestration.	VM	OpenStack/Kubernetes
PiaFFE [74]	vNF orchestration	vNF Offloading to eNF, improved throughput and eNF placement.	VM	Servers with SmartNIC

4.1.1. Eden Framework

The Eden framework was proposed by Ballani et al. [70] as a framework suitable for virtual network function provisioning on end-user devices. Eden leverages the concept of data-plane programmability by implementing NFs on end-user devices written in the F# language, which is a high-level programming language. Eden comprises three functional components that work together to handle packet processing tasks: the centralised logical controller, which is used for service orchestration and for providing a global view of the infrastructure. The second component is the *Stage*, which is simply a name for kernel modules, libraries, or applications on hosts that are used for the classification of packets before being sent to the Enclave (the third component in Eden).

The Enclave is responsible for handling the functionalities of a programmable data plane, which can be implemented on NICs, FPGAs, hypervisors, or operating systems. Eden maintains match-action rules in the Enclave, with traffic association carried out on the host device by the stage component of the framework. The Enclave is also responsible for interpreting the bytecode, which is obtained from compiling action functions. Eden still leaves the open question about where to best deploy the network functions, that is, either on the user OS or on the programmable NIC. In Section 5, we shed some light on this open challenge by arguing for a hybrid implementation framework for fast packet processing and efficient service chain creation in next-generation networks.

4.1.2. E2 Framework

The E2 framework presented by Palkar et al. [71] is designed for the management of NFV applications and resource allocation, which is achieved without necessarily knowing

the low-level implementation of the applications. The target environments for the E2 system are hardware commodity servers and switches in high-performance network environments, which are typically found in today's central office locations. E2 implements a manager, which is responsible for orchestrating communication between the SDN controller and a cluster of servers.

The E2 framework also manages the placement of NFs using the proposed algorithm on available servers by monitoring the available resources and efficiently placing NFs to avoid unnecessary system overheads. Network operators can define their policies using *pipelets*, which state the steps involved in processing traffic from a specific class. A directed acyclic graph (DAG) is a key component of the E2 *pipelet*, which defines how a class of traffic is processed by the E2 NFs, with nodes representing physical switch ports or NFs.

The configuration of network functions is done by providing the following inputs: (i) an API exported by E2 for leveraging optimisation options, (ii) a method for attribute association, that is, for per-packet and port metadata; (iii) information on the scalability of the application, that is, whether it can scale across multiple cores or multiple servers, which gives the E2 framework an idea of how to handle situations of traffic overload, (iv) a method of splitting traffic across multiple NF instances by considering the constraints of the target environment; and (v) information on the NF processing capacity in terms of traffic rate, which helps with placement decisions. E2 provides service providers with the flexibility of declaring their policies without prior knowledge of the underlying network function or infrastructure.

4.1.3. Pishahang Framework

Pishahang is a multi-domain service orchestration framework for SFCs, proposed by Kouchaksaraei et al. [73,75] (with the introduction of dynamic service chaining), which combines container-based and VM-based vNFs to create a chain of network services. The framework was implemented across the OpenStack and Kurbenetes domains. In terms of service description, Pishahang uses two descriptors: the first for describing information considered to be *high level*, such as service chaining, microservices, and vNFs, while the second descriptor focuses on a more fine-grained description of vNFs, such as the required resources for running the vNFs. Pishahang is a framework that has been built to chain services across heterogeneous domains.

Pishahang used the SONATA MANO https://www.sonata-nfv.eu/ (accessed on 15 December 2021) framework, which supports the addition of new functionalities following a microservice-based architecture. Service graphs are translated by the SDN adaptor, which is sent to the controller, converted to forwarding rules, and installed on switches. To validate the features of Pishahang in chaining services across multiple domains, VM-based and container-based forwarders were chained to create an SFC with ICMP packets sent end-to-end. The use of containers for the deployment of vNFs is yet to be fully developed because of reasons such as lack of functional isolation, and security.

4.1.4. SNF Framework

SNF is a SFC framework proposed by Katsikas et al. [16], which synthesises SFCs with the main goal of performance optimisation by eliminating redundancy in packet processing across the service chain. Typical causes of redundancy are eliminated by the SNF framework by (i) creating a logical processing entity to handle all chain-wide operations on received packets, rather than handling network functions as separate processes, (ii) discarding packets that need to be discarded very early in the service chain; (iii) reducing multiple read operations by collecting read operations and constructing classes of traffic as a directed acyclic graph, which is synthesised into a classifier, and (iv) reducing the number of write operations by modifying traffic classes in a single operation.

The concept of set theory and graphs is employed for traffic classification to achieve the synthesis of similar network functions to improve the overall performance of the service chain. Another key feature of the SNF framework is the management of states

across multiple network functions in a service chain, which enables the synthesis of stateful service chains. At any given time, there is a processing core that actively classifies all the received frames into the required traffic class units. The ingress traffic is hashed using RSS, which helps to serve bi-directional flows by the same processor and re-writer [16]. Figure 9 depicts the SNF framework on a device with two network cards, where each NIC is tied to a CPU core.

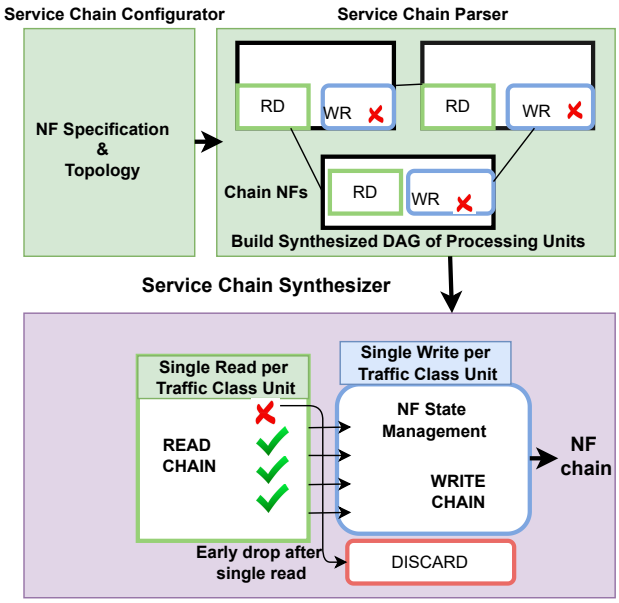

Figure 9. SNF Framework.

4.1.5. Open Baton Framework

The Open Baton framework was proposed by Carella et al. [72] as a framework for NFV service orchestration and management. The framework has the OpenStack cloud infrastructure as its underlying platform, which is compliant with the ETSI MANO architectural framework. The key component implemented by Open Baton is the multi-site service orchestration feature in heterogeneous network environments. Open Baton provides a management component for handling the life cycle of network functions, including the use of the JUJU http://openbaton.github.io/documentation/vnfm-juju/ (accessed on 20 December 2021) virtual network manager for interoperability between diverse network functions.

The framework provides support for diverse VIMS, which means there is no need to rewrite the components of the logic that is responsible for service orchestration. To speed up the NF instantiation time, drivers are provided for VIMs and VNFMs that support the deployment of containerized network functions. Scaling of network functions can be handled at runtime using the auto-scaling component https://github.com/openbaton/autoscaling-engine (accessed on 1 January 2022).

Open Baton uses Zabbix http://openbaton.github.io/documentation/zabbix-plugin/ (accessed on 1 January 2022) to monitor network activities and network function status. The framework handles resource allocation by leveraging the concept of network slicing using SDN; thus, the extensibility and interoperability of Open Baton makes it ideal as an orchestration framework for heterogeneous network functions. A fault management module, together with a management dashboard, makes Open Baton a complete solution for heterogeneous NFV orchestration.

4.1.6. Metron Framework

Metron is an NFV framework proposed by Katsikas et al. [17], which achieves high utilisation of commodity servers and underlying network resources. Metrons can offload some packet processing tasks to the underlying network infrastructure and achieve low inter-core communication using tag-based hardware dispatching for processing packets.

The reduction of inter-core transfers implemented in Metron gives it the capability to process packets at the speed of the L1 cache. Metron performs stateless packet processing and classification by leveraging the OpenFlow and P4 protocols.

The problem of having a mismatch between the server and network architecture is also addressed by tagging packets to be dispatched and switched in the service chain, which is controlled by the implementation of the ONOS SDN controller. Placement decisions of synthesised packet processing graphs are carried out accurately and at a low cost by obtaining the network state. A load-balancing scheme was introduced for servers and CPU cores [17]. Figure 10 presents an overview of the Metron architecture, with an example network function chain and execution steps consisting of a firewall and DPI NF.

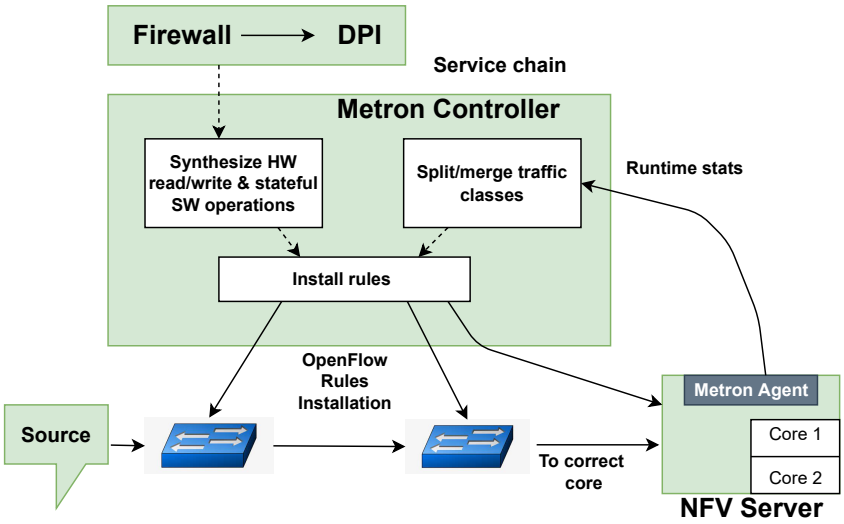

Figure 10. Metron framework overview.

4.1.7. CN-SFC Framework

Dab et al. [76] presented an SDNLess SFC microservice architecture for Cloud-Native NFV, a framework for cloud-native SFC creation which uses an advanced version of network service mesh (NSM) https://networkservicemesh.io/ (accessed on 27 December 2021) and Kubernetes (https://kubernetes.io/) (accessed on 2 January 2022) for chaining cloud-native elements to form a service chain. They considered the use of SFCs in micro-service architectures and addressed the shortcomings of the NSM architecture, such as its inability to handle L2 or L3 traffic and the lack of support for advanced routing algorithms.

The cloud-native framework, CN-SFC, was also used to achieve traffic steering by efficiently load-balancing traffic across the Cloud-native functions (CNFs). An approach was introduced, which uses the weighted round robin algorithm by maintaining the weights of Kubernetes pods and distributing traffic based on pods with high weight values. TS-CNF, which is a traffic steering problem, was modelled and solved as an integer programming (IP) problem.

In terms of the evaluation of CN-SFC, a Kubernetes cluster was used, which consists of an NSM control plane, a master node, and two worker nodes. The Docker and the Kind https://kind.sigs.k8s.io/ (accessed on 28 December 2021) framework were used to run the cluster. The proposed network-aware traffic steering scheme (NA-TS) was evaluated by varying the number of flows that need access to the network service in the cluster. The number of replicas that were used for deploying the VPN and firewall services were also varied, while UDP and ICMP (to evaluate reachability) packets were generated and evaluated (jitter and packet loss in the service chain) between the two assigned pods and those deployed in the cluster. Additional metrics such as the latency of the network and service instantiation time were also evaluated.

4.1.8. Slick Framework

The Slick framework was proposed by Anwer et al. [15], which allows for the programming of network functions with a control program that describes how packets can be processed in a specific traffic set. Traffic flow is specified by applications that provide information on the elements that need to be traversed by the packets in the network. The controller deploys the selected processing elements and deploys them on the machine. The task of resource management in the framework is handled by the slick-run-time, which ensures that the links and processing elements are not overloaded with packets beyond capacity.

Network functionalities are placed as software elements, which can be installed dynamically or at the initialisation time. Events are sent to the Slick controller by the modular processing elements; thus, Slick takes care of the placement of network functions and the steering of traffic between packet processing elements. In terms of evaluation, the proposed framework was evaluated using a Mininet SDN network emulator with various network topologies. The controller was run on a separate VM, while the emulator was run on another VM with 60 SDN switches. Both VMs had eight CPU cores. One of the goals of Slick, which was achieved, is to maximise bandwidth utilisation between various network functions.

Slick allows the network programmer to use a high-level language to describe a module that handles the steering of traffic, as well as the placement of lightweight functions in arbitrary locations along the service chain. Slick does not consider other environments, such as the edge of the network; thus, their implementation is restricted to data-center environments.

4.1.9. Openbox Framework

The OpenBox framework proposed by Bremler et al. [14] is a framework for the management, development, and deployment of vNFs. OpenBox abstracts network functions as packet processing graphs, which represent the behaviour of typical network functions, such as a firewall, DPI, and NAT. Processing graphs are implemented using the elements of the click router framework [77], specifically for firewall, web cache, load balancer, and IPS.

To improve performance, OpenBox introduced a graph merging algorithm to merge the abstracted NF processing graphs, which reduces per-packet latency by minimising the number of processing blocks that need to be traversed by packets. The graph merging process starts by normalising the graphs into trees to avoid the convergence of paths, and the resulting trees are concatenated in the right order of packet processing by the network functions (to ensure chain-wide correctness of the processing pipeline).

There are three major components that make up the OpenBox framework: OpenBox Applications, OpenBox Controller (OBC), OpenBox Service Instances, and OBIs, which constitute the OpenBox data-plane. Packet processing graphs are sent down to the OBI data plane (which can be implemented in hardware or software) from the controller, which in turn receives information about the packet processing capabilities of the OBIs. The communication between the OpenBox controller and the data plane is handled by the OpenBox protocol, which defines packet processing blocks for building vNFs.

The controller is used by OpenBox to achieve multi-tenancy, smart NF placement, and NF scaling, in addition to steering traffic to the right vNF. This component of the OpenBox framework provides network application developers with a layer of abstraction for creating applications that have a specific packet processing graph and logic. The OpenBox framework was evaluated using service chain configurations and pipelined network function scenarios, with both scenarios yielding much higher throughput when compared with scenarios without OpenBox. To cater to the resilience of network functions, the authors merged multiple NFs together to create a single processing pipeline, which provides higher throughput to one of the network functions at off-peak times of the second NF in the merged processing graph, yielding a throughput that is 20% better than the naive merge approach.

4.1.10. Piaffe Framework

PIaFFE, a Place-as-you-go in-network framework for flexible embedding of vNFs, is a placement framework proposed by Mafioletti et al. [74], which achieves multilevel chaining and placement of vNFs implemented on SmartNICs. The overall aim of the PIaFFE framework is to maximise throughput and achieve minimum latency by embedding network functions (eNFs) on in-network processors. These embedded network function implementations reduce server CPU utilisation on end-hosts, increase throughput to line rate speeds, and reduce latency.

A PoC which was implemented on two physical servers was achieved by chaining three network functions for authentication, IDS, and firewall functionality. PIaFFE reduces host load by performing full or partial vNF offloading while consolidating multilevel chaining. PIaFFE made use of the P4 programming language to steer traffic to vNFs or eNFs, depending on what function has been embedded as an eNF. This decision is determined by the P4 hash table or bloom filter, which has been implemented using a P4 Data Structure (P4DS).

4.2. Performance Tuning Frameworks

Improving the performance of virtual network functions, either as standalone functions or as part of a service function chain, helps network operators reduce service instantiation costs [78], and the optimal use of available resources. In this subsection, we present frameworks that focus on optimising the performance of network functions in a service chain. Most frameworks that are designed to optimise performance devise mechanisms that can achieve goals such as reducing network function deployment and provisioning time, maximising the throughput of network applications and reducing latency (Table 4). We would like to note that there are several efforts in the literature that focus on optimising network functions using diverse technologies and implementation methodologies. We present the state-of-the-art equivalence classes of frameworks in this category by explaining the design, technology, methodology, and results obtained by each framework.

4.2.1. Coco Framework

CoCo is an NFV framework proposed by Meng et al. [79], which was designed for the deployment of modularized service function chains (MSFCs). One of the major goals of the proposal is the consolidation of processing elements collocated on a VM and handling the placement of the modularized SFC to minimise packet transfer overhead between VMs. The authors designed a placement scheme that selects the right SFC elements for consolidation using performance and resource-aware placement. Fairness is achieved between several NFs tied to a single CPU core using a run-time scheduler implemented in CoCo.

In terms of scalability, CoCo can utilise a push-aside scheme specifically designed to handle reduction in performance, which might arise due to scaling elements. Unlike most existing NF scalability approaches that start up a VM when there is a need to scale, which leads to more overhead in terms of latency, the push aside algorithm reduces the need for inter-VM hop creation. Rather than creating a new replica (as used by traditional vNF scalability solutions), the CoCo framework adds more resources to elements that are overloaded.

In terms of performance evaluation, the efficiency of resource utilisation and the reduction in the cost of transferring packets across the service chain are the two major benefits of the CoCo framework. Throughput and CPU utilisation were measured with CoCo implemented using Docker containers to allow for the consolidation of processing elements; thus, Open vSwitch was employed as the virtual switch for VM-VM communication. The results show that CoCo can improve performance by approximately 45.5% and a 2.46X reduction in packet transfer overhead.

Table 4. Performance tuning frameworks.

Framework	Problem	Solution(s)	Processing	Testbed/PoC
ClickOS [12]	Middlebox optimization	Optimized packet processing and Zen-based VM optimization.	Click	Xen-based hypervisor
OpenNetVM [13]	VNF performance optimization	DPDK-enabled vNFs, containerised NFs, kernel by-pass packet processing and line-rate packet processing.	Containers	DPDK-enabled hosts and containers
NetBricks [18]	NF performance optimization	Zero-copy isolation and memory level isolation.	VM	Physical servers
GNF [19]	Edge vNF deployments	Lightweight vNFs and edge service chains	CN	Host with CNs
CoCo Framework [79]	Modular SFC	Optimization of modular SFCs, effective resource allocation, resource-aware placement, vNF fault tolerance and optimized run-time scheduler.	Containers	Commodity servers
DeepNFV [80]	NFV SFC performance	Deep learning for QoS and traffic optimization, lightweight containerised vNFs and Edge NF deployments.	Containers	GNF framework
NetFate [81]	Edge NFV deployments	Active NF/VM migration and open PaaS platform.	VM	Commodity servers and virtual switches
HyperNF [82]	NFV SFC performance	Reduced I/O sync overheads, hypercall-based I/O in VM context and vNF scalability.	VM	VALE switches and VMs
PhantomSFC [83]	Resource optimization	Service/Control decoupling, reduced latency and improved throughput.	VMs	VM/DPDK
MVMP [84]	NFV SFC performance	Improved throughput, lightweight vNFs, vNF fault tolerance and vNF replication.	Containers	DPDK, Container, Virtual Switch
MicroNF [85]	NFV/SFC performance	Optimal NF placement, reduced inter-NF latency, vNF performance optimization and scalability and fair scheduling of vNFs.	Containers	Docker and VMs

4.2.2. Deepnfv Framework

DeepNFV is a lightweight NFV framework proposed by Li et al. [80], which was designed specifically for edge network deployments, with the aim of minimising the packet processing tasks at the core of the network by offloading to edge network functions. DeepNFV is built on the GNF framework [19], which uses lightweight docker containers to build network functions for the edge. The key components of the proposed framework are the deep learning models employed and the infrastructure layer, which handles the interaction between network links and devices.

DeepNFV uses deep learning to enhance tasks such as the optimisation of QoS parameters, classification of traffic, and analysis of network links. Similar to the GNF framework, DeepNFV was built to support the idea of moving network processing elements as close to the data source as possible (edge computing). As a use case for the DeepNFV framework, network traffic analysis functionality was considered by the authors by generating basic images from network traffic. A traffic analysis-containerized network function was used for the analysis and classification of images using deep learning models.

To demonstrate the traffic analysis use case, the DeepNFV framework starts by splitting the received traffic into discrete components, which are stored as PCAP files, and the second step involves the modification of the packet headers to trim the header length or remove unimportant fields from the packets. The modified PCAP files are *cleaned* to remove duplicates before being converted into image data. The resulting images are processed by

the CNN model and sent to the next network function in the chain for further action(s). The ability of the framework to classify images and the performance of the network functions at the edge of the network were evaluated, and improved performance in terms of precision and efficiency was recorded.

4.2.3. Microf Framework

The MicroNF framework was proposed by Meng et al. [85] as a framework for the deployment of modularized service chains, using a centralised controller for service chain graph reconstruction and redundant NF reuse. Service providers describe the MSFC to be deployed by clearly defining how the elements are interconnected. This is followed by processing the MSFC using graph reconstruction and the identification of any dependencies by elements, with the aim of reusing elements where possible. The reordered MSFC is optimally placed with the goal of reducing the latency between processing elements.

The major goals of the framework are to (i) efficiently reuse elements that have similar configurations in the processing pipeline by first addressing the problem of dependency between different elements, (ii) solve the problem of VM to VM connection using a virtual switch in an optimal fashion, (iii) shorten the service chain length, and reduce packet processing costs, where necessary.

In terms of the scalability of NFs, MicroNF implements run-time scaling algorithms, which ensure minimal inter-NF latency along the service chain. The problem of selecting processing elements that are ideal for consolidation is also handled by the MicroNF framework, in addition to a placement algorithm that prioritises high performance. The speed of packet processing between diverse elements is also considered by the proposed resource scheduler, which ensures that the workload is efficiently shared among available processing elements [85].

4.2.4. Netbricks Framework

NetBricks is an NFV framework proposed by Panda et al. [18], which offers a platform for building and running virtual network functions that provide software isolation between NFs. The NetBricks framework differs from other approaches by (1) limiting the set of processing modules to core functionalities, which helps to reduce the number of modules that network application developers must deal with, and (2) allowing the customization of modules using user-defined functions, which makes the modules more flexible and optimised for better network function(s) performance.

NetBricks eliminates overheads resulting from context-switching by enforcing memory-level isolation in software and reducing I/O related overheads by introducing zero-copy software isolation [18]. Using zero-copy isolation, the cost of packet I/O is greatly reduced by NetBricks, which means that chains of network functions can be run as a single process.

NetBricks provides a major distinction in providing fault and memory isolation for NF implementations by utilising operators designed for parsing, de-parsing, transforming, and filtering packets. In addition to packet operators, the framework also provides abstractions for processing byte-streams, abstractions for control flow, and for state and scheduled events. To evaluate the performance of the NetBricks framework, two example network functions were used: the first is a simple network function that decrements the TTL of a packet and discards any packet that has a TTL of 0, and the second is a stripped down implementation of the Maglev load balancer [86], which splits ingress traffic among servers and also provides failure recovery for back-end servers.

The measurements evaluated include simple NF overheads, array bound overheads, and how *general* the NetBricks programming abstractions can be. For the latter part of the evaluations, that is, the programming abstractions, five network functions were implemented: NAT, firewall, Maglev load balancer, and a Snort-like NF that performs signature matching on ingress packets. Improved performance was observed for scenarios where (1) CPU cores and chain lengths were varied, (2) the load was varied with respect

to CPU cycles and chain length, and (3) throughput measurements for single network functions with a variable number of CPU cycles for multiple isolation approaches.

4.2.5. Hypernf Framework

HyperNF is a high-performance NFV platform proposed by Yasukata et al. [82], which aims to properly utilise commodity server resources while scaling the number of network functions hosted by servers. The problem space addressed by HyperNF includes resource allocation, efficient utilisation, and high throughput when using everyday commodity servers to deploy virtual network functions.

The proposed framework is aimed at large NF deployments, where utilisation is maximised for better throughput. The use of hypervisor-based I/O is employed, which helps reduce synchronisation overhead. HyperNF was designed using three core design objectives: (i) CPU cores are not reserved entirely for virtual I/O operations, thus providing high flexibility in terms of utilisation, (ii) proper accountability for virtual I/O tasks on respective VMs, thus offering a cohesive resource allocation strategy, and (iii) VM switches should not be used for packet switching; instead, the data path of software switches is exported to the hyper-visor for the purpose of forwarding and switching of packets.

HyperNF was evaluated for scenarios involving a baseline setup using the VALE [86] switch for inter-VM communication, with each VM tied to a single CPU core, and second, a scenario that consolidates network functions in a shared CPU environment by varying the number of VMs (CPU cores are shared among the firewall VMs deployed using a round-robin scheme). Both scenarios outperformed the split and merge schemes compared with HyperNF. Other tests carried out include resource allocation, NFV throughput, and SFC chain composition. A chain of 50 NFs can achieve a delay as low as 2 ms, which makes the framework ideal for SFC deployments.

4.2.6. Netfate Framework

NetFate was proposed by Lombardo et al. [81] as a framework that supports the deployment of network functions at the network edge and data centre infrastructure. The main elements of the NetFate framework are simply the clients, which receive or generate packets, and the CPE nodes, which hosts the network functions for clients to connect to the infrastructure.

The orchestrator contains an SDN controller for handling communication with OpenFlow switches, an NFV coordinator for handling VM life-cycle and hypervisor-VM communication, and an orchestration engine which collects statistics about available devices, connected clients, and network services. Each time ingress packets are received, the orchestrator (i) takes a decision on which NFVI can host the NF based on defined SLA, (ii) carries out the migration or instantiation of VMs for hosting the NFs in (i), (iii) creates a virtual service path for connecting VMs that host the NFs, (iv) forward ingress flows based on defined routing policies, and (v) terminating unused VMs, thus making resources available.

The Proof of Concept employed for the evaluation of NetFate comprises client devices and nodes that represent network access points, a controller, and an orchestrator which also authenticates and authorises users. This was made possible by the implementation of two firewall network functions at CPE nodes, where migration efficiency is measured while moving from one CPE node to the other. The NetFate framework is ideal for customer premise equipment network function deployments; thus, its performance in terms of provider equipment implementation is yet to be evaluated.

4.2.7. Clickos Framework

ClickOS was proposed by Kohler et al. [12], which uses the *elements* from the Click software router [77] to achieve lightweight middlebox packet processing. This runs on Linux VMs and an XEN-based https://xenproject.org/users/virtualization/ (accessed on 4 December 2021) optimised platform. To achieve domain isolation, each click middlebox is run on a separate Linux VM, which provides memory isolation. ClickOS achieves high

performance in terms of network I/O by making the following changes to the Xen network pipe: (i) replacing the OvS backend switch, which makes it easier to map VM memory, (ii) moving the netback driver to the control plane, which serves as a communication medium with the netfront driver, and (iii) modifying the netfront driver of the VM to allow the mapping of ring buffers. The framework presented some useful modifications to the Xen backend and frontend modules to achieve faster data transmission rates.

The evaluation results show that ClickOS speeds up networking for Xen-based VMs by applying several well-known optimisation approaches, such as removing unnecessary data paths and batching of processes. The performance of the ClickOS switch was measured, in addition to metrics such as boot time, memory footprint, throughput, delay, chaining, scalability, and middle-box state insertion. An increased throughput from 8 Kp/s to 344 Kp/s was achieved by changing the driver settings, receiving grants for buffers at the initialisation time, and re-use the buffers for all packets. The VALE switch, which is an in-kernel virtual switch in Linux that allows for scalability in terms of the number of ports and throughput, was replaced by Open vSwitch.

4.2.8. Opennetvm Framework

The OpenNetVM framework was developed by Zhang et al. [13] as a framework that is ideal for high-performance vNF deployments using the Intel Dataplane Development Kit (DPDK) https://www.dpdk.org/ (accessed on 8 December 2021) and Docker Containers https://www.docker.com/ (accessed on 17 December 2021). To achieve high-speed packet I/O transfer, OpenNetVM implements zero-copy to reduce the I/O overhead associated with copying packets from the NIC for processing in the user space. Packets are DMA'd directly from the NIC to a shared memory space, which is accessible to DPDK-based network functions supported by the framework.

As depicted in Figure 11, a shared memory space is created by the manager, which stores the metadata information, list of service chains, and flow tables. Dependencies are encapsulated in Docker containers that host the network functions; thus, packet transfer between NFs is handled by the TX and RX threads that carry useful descriptors. High-speed packet processing is also made possible by utilising the DPDK poll-mode driver rather than interrupts.

In terms of NF-to-NF communication, OpenNetVM uses a centralised logical controller, which communicates using the OpenFlow protocol to orchestrate NF activities. The network function manager is responsible for managing memory and the NF life-cycle by handling inter-NF communication and sending keepalive messages. The transfer of packets between the NICs and NFs is also handled by the manager using the TX and RX threads. Network functions are either implemented using DPDK or as a user-space container process, which are tied to specific CPU cores in both scenarios. The framework was evaluated using metrics such as the scalability of multiple ports, overflow of the flow director, performance of service chains, and the flexibility of the framework can steer packets, which yields a much better performance when compared with ClickOS, especially in terms of throughput with variable chain length.

4.2.9. Phantomsfc Framework

PhantomSFC was proposed by Castanho et al. [83], which is an SFC framework aimed at decoupling the underlying network from the service plane. The design was implemented to be transport-independent, network agnostic, elastic, and maintains a small footprint by considering end-to-end throughput and latency. The PhantomSFC framework is based on the IETF SFC reference architecture presented by [37], which comprises a classifier, an SFF, SF, and a proxy component. NSH, which was standardised by [87], was used as the encapsulation protocol for SFCs in PhantomSFC.

Components such as proxies, forwarders, and classifiers are deployed as vNFs, and a centralised SDN controller is employed for the realisation of chain configuration and instantiation. Tasks involving chain configuration, such as creating a new chain and

removing and modifying configuration rules in proxies, classifiers, and forwarders are carried out by the logically centralised controller. Using PhantomSFC, resources can be scaled by SPs based on service demands; thus, the PoC evaluation of PhantomSFC achieved improvements in throughput, jitter, and latency, using the DPDK application.

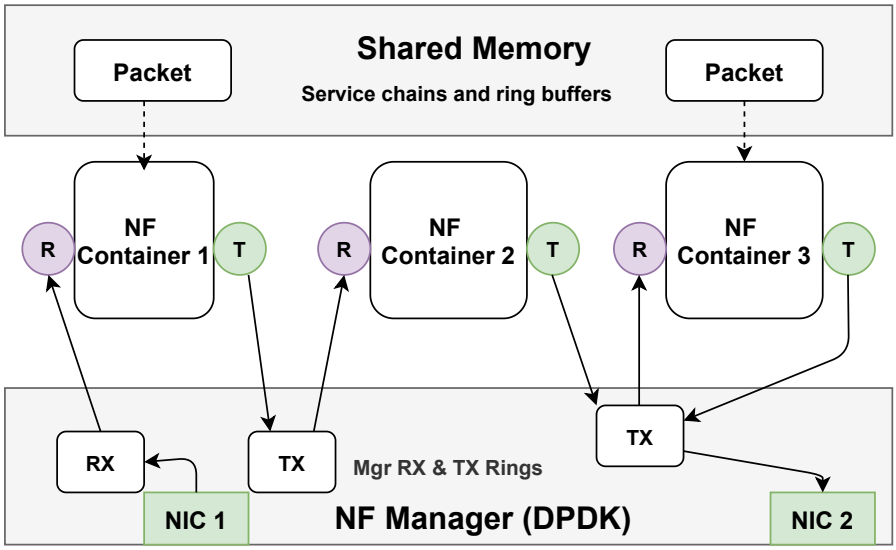

Figure 11. OpenNetVM framework.

4.2.10. GNF Framework

The Glasgow Network Functions (GNF) framework was proposed by Cziva et al. [19] as an open and lightweight implementation framework for network functions in OpenFlow network environments. GNF, which is a container-based framework, achieves low overhead in terms of performance, high NF reuse, and fast deployment speed, when compared to most NFV deployments. To achieve the routing of traffic in a typical NFV scenario, policies can be implemented by adding entries or by adding a middle-box in the path of the traffic. These two approaches have drawbacks that the GNF framework attempts to eliminate.

GNF achieves dynamic placement of network functions by simply rerouting the traffic to the server with the requested NF, which allows service providers to utilise the same hosts when handling network and compute functions, thus minimising the overall infrastructural costs. Using the OpenFlow protocol, GNF can match ingress packets to the match-action table before routing packets to the specified destination.

For ease of network function deployment and management, GNF provides a user interface for global control and view of the network, a manager, GLANF router, and agent. The life cycle of network functions is handled by the GLANF Manager, which makes use of the OpenDaylight SDN controller for performing tasks such as creating, starting, stopping, and deleting primitives (Figure 12). Tasks such as resource allocation are also handled by the manager, which allocates network functions to hosts that have available resources.

To evaluate and demonstrate the performance of the GNF framework, six network functions were deployed: a wire, which routes packets from its ingress to egress ports, an HTTP filter, traffic control, a load balancer, intrusion detection, and a firewall, which is based on iptables. Figure 12 depicts the GNF framework, with packets sent from VM1 to VM2, which are sent via the network functions on the GLANF system and to target hosts. The performance of GNF was also compared to ClickOS [12], which shows a significant improvement from 3.6 Gb/s to 13.8 Gb/s packet processing speeds, this also holds true as the number of chained containerised vNFs are increased.

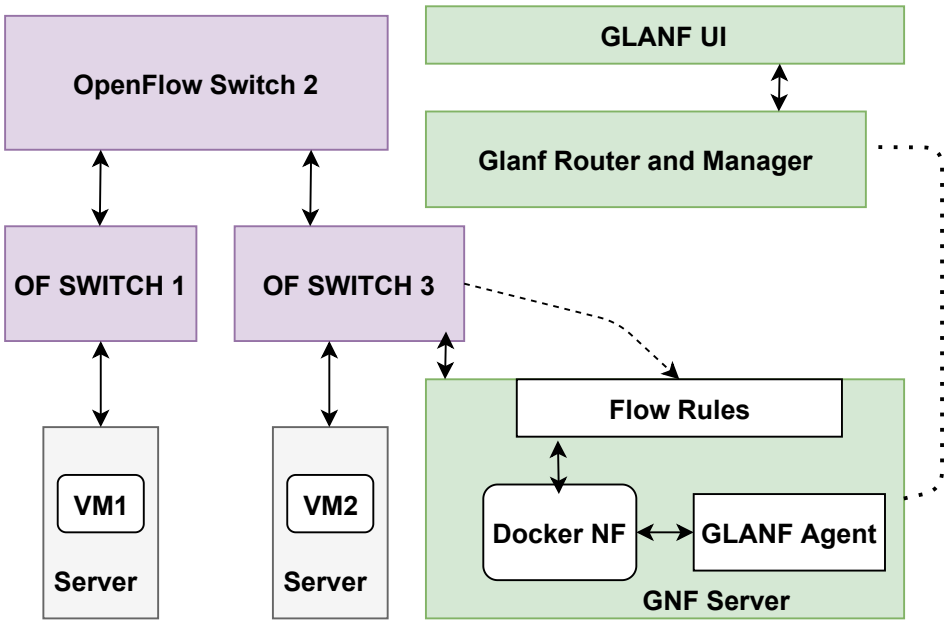

Figure 12. GNF architecture.

4.2.11. MVMP Framework

The NFV platform proposed by Zheng et al. [84], i.e., the Multiple Virtual Middlebox Platform (MVMP) is a high performance framework that has been built using the Intel DPDK platform and Docker containers. The three major components of the proposed MVMP framework are (i) abstracted virtual devices, (ii) a control plane, and (iii) a shared memory space. Packet processing by NFs is achieved by an abstraction layer that supports the deployment of multiple NFs on a single hypervisor. Network functions are run in user space as processes, which makes them lightweight, thus requiring fewer resources for packet processing. Packets are polled directly from the NIC using the DPDK poll mode driver and sent to several network functions, which also adds to the fast packet processing speed of the proposed framework.

In terms of implementation and evaluation, network functions were implemented and chained together, and the service chain performance was evaluated and compared with the OpenNetVm [13] framework, which yields 3x better throughput as the number of network functions is increased in the chain, with an overhead of approximately 4% with regard to network function isolation.

4.3. Resilience and Fault Recovery Frameworks

The resilience of virtual network functions to link, node, and chain-wide failures has been addressed using diverse technologies and methodologies in the literature. The state of all network functions (active and standby) in a chain is vital when creating a resilience mechanism; thus, building a fault-tolerant middlebox and service chain becomes imperative [88]. Different frameworks make use of various mechanisms to detect faults, fix them, and resume normal packet processing operations with as little downtime as possible [89]. This section presents equivalence classes of frameworks that focus on the resilience or survivability of virtual network functions in a service chain (Table 5), by explaining the design choices and implementation technology used.

Table 5. Resilience and fault recovery frameworks.

Framework	Problem	Solution(s)	Processing	Testbed/PoC
Medhat et al. [48]	SFC Resilience	Runtime traffic reroute for fail-over and fault recovery.	VM	OpenStack/OpenDalight
CN-SFC [76]	Traffic Steering	NSM and network-aware steering.	Containers	Kubernetes
FTC [88]	SFC resilience	Chain-wide vNF fault tolerance, vNF state piggybacking and resource management.	Click	Server cluster
REINFORCE [90]	SFC failures	Single and multiple nodes failure recovery, remote and local redundancy, link and node failure detection.	Containers	Physical servers
Hmaity et al. [91]	SFC Resilience and placement	Single Link/Node failures.	VM	ILP Models
FtVNF [92]	NF fault tolerance	Slave and Master vNFs deployment, vNF failure recovery and fault tolerance.	Click	Commodity servers
Nguyen et al. [93]	SFC Resilience	Controller-independent HA scheme.	VM	OpenFlow/OpenStack

4.3.1. Reinforce Framework

Kulkarni et al. [90] proposed REINFORCE, a framework for achieving resiliency of DPDK-based NFs, which provides the check-pointing of applications that reduce the state of network functions to be replicated. REINFORCE provides failure recovery of network functions across the entire chain, with the detection of node and link failures within the shortest possible time. Packet processing overhead is minimised by the separation of network function behaviour into deterministic and non-deterministic, thus committing to check-pointing the states of standby NFs in non-deterministic scenarios.

REINFORCE emphasises stateful network functions, which maintain the state of connections either globally or per-flow. The characterisation of state information enables the framework to decide whether flow updates are deterministic, which helps with the synchronisation of NFs that operate in a particular chain. The use of lazy check-pointing of the NF state and the replay of packets is used by REINFORCE to speed up the process of recovering from failures. Figure 13 depicts the architecture of the framework in which a chain-wide symmetry is maintained by nodes. NFV nodes can host multiple network functions, which can either be part of a service chain or a complete chain in a single NFV node. The NF manager is able to access the shared memory pool, while the process of time-stamping ingress packets is carried out at the beginning of the chain, and sent to the next NFV after logging.

4.3.2. FTC Framework

Ghaznavi et al. [88] presented a framework for fault-tolerant chaining (FTC). FTC uses a different approach, which opposes existing solutions where middle-box snapshots are taken for the purpose of replicating state, or approaches where the state of middle-boxes is stored in a fault-tolerant data store. The design requirements of the FTC framework are (i) correctness of middlebox recovery, (ii) quick recovery from failures with low processing overhead, and (iii) efficient use of servers hosting middleboxes. Middlebox state information is added to the packets as they traverse the SFC chain, which is replicated in the host servers. The deployment of fault-tolerant chains is handled by the ONOS controller, which serves as a centralised orchestrator for the management of NF and chain life cycles.

To achieve fault tolerance, FTC makes use of replicas, which comprise data and control plane modules for interacting with the orchestrator. New threads are spawned by the control module in fail-over scenarios [88]. To optimise the amount of memory used for service replication, updates that have been added to the standby middleboxes are removed. FTC middleboxes are built using Click [77] elements that interact with the ONOS controller. Parameters such as replication factor, time required for failure recovery, throughput, and

latency were all measured while varying the length of the service chains deployed first in a cluster of 12 servers and second on distributed servers on the cloud. The FTC was compared with FTMB [94] and NF, which is a baseline framework designed with no fault tolerance. FTC produces a much higher throughput (with an increase in service chain size) with a chain-wide overhead that is less than that obtained with FTMB.

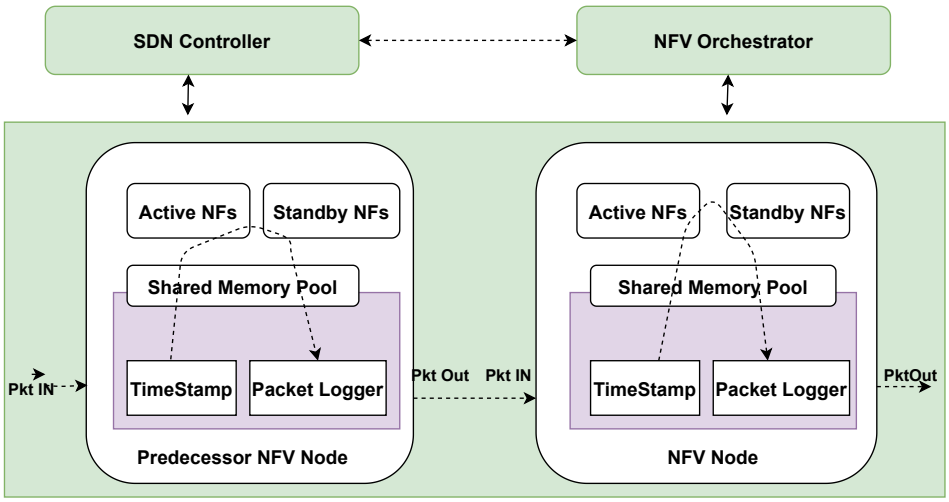

Figure 13. REINFORCE framework.

4.3.3. Hmaity et al.

Hmaity et al. [91] presented a solution that uses integer linear programming to solve the problem of link and node failures in service chains. They modelled the physical infrastructure using a directed network graph, with representations for physical devices with the ability to host virtual network functions and links. The latencies of physical links and packet processing latency by vNFs are also represented in the model of the physical network topology. The other components that were modelled are the service chains and the vNFs, which are considered as abstract components that can process ingress packets before forwarding to the next vNF or host device.

The constraints considered for the proposed model are (i) node capacity and link latency, which capture the current state of a node, the capacity of the link in use, the latency and the highest number of virtual machines that can be hosted by a particular node; (ii) constraints related to routing, such as the location of virtual links in the physical network, to ensure that routing paths are made available on the right physical node; and (iii) constraints related to the placement of vNFs on physical devices, by ensuring that active and standby vNFs are not collocated on the same node.

The proposed models were evaluated using two example service chains, that is, a chain consisting of a web service and another service chain that models online gaming, where the impacts on node capacity and latency were considered. To solve the formulated ILP model, a bandwidth of 100 kbit/s was set for the online gaming service chain, while 50 kbit/s was set for the web service scenario. Latencies of 500 ms and 60 ms were set for the service chains. The proposed and solved models showed that achieving SFC resilience requires additional (redundant) nodes of approximately 107%. Although the proposed solution serves as a good mechanism that can handle link and node failures in service chains, it lacks the ability to provide shared protection against failures.

4.3.4. Resilient SFCs—Medhat et al.

Medhat et al. [48] proposed an OpenStack and OpenDaylight-based environment to deploy and orchestrate resilient SFCs in cloud environments. The proposed framework follows the ETSI NFV model, which is capable of traffic rerouting, in the case of faults occurring at runtime. The authors proposed an extension to the ETSI NFV framework for

SFC orchestration and management, with the implementation of a service chain consisting of two firewall network functions in standby and active modes.

They used the Open Baton framework [72] as the NFVO, which uses a messaging queue to communicate with the SFC orchestrator (the OpenDaylight SDN Controller), and the Zabbix network monitoring tool. Service function failure is simulated by abruptly terminating the process running the NF, then the Open Baton Fault Management System (FMS) switches to the standby network function, and the failed NF is recovered by the Orchestrator.

4.3.5. OpenFlow Fault Recovery

The framework proposed by Nguyen et al. [93] used the OpenFlow group table to provide a quick fault recovery and fast fail-over scheme. Their proposal eliminates the use of a centralised logical controller by exploiting the use of the OpenFlow group table for managing service function chains. Fault recovery and detection were implemented in local OpenFlow switches, which utilise an OF group table.

The need to contact the controller or NFV-MANO in the event of a failure is eliminated, which saves time on fault notification and recovery. The framework was tested using OpenStack and OF, with two vFirewall functions deployed for redundancy, which showed a reduction in SFC packet loss and an improvement in link throughput, as well as quick failure recovery. We briefly explain the aspects of SFC resilience that require further attention from the research community in Section 5.4.

5. State-of-the-Art and Open Challenges

The challenges that the classified frameworks in Section 4 try to solve and the approach used are presented in Tables 3–5. This captures the proposed solutions, the packet processing element(s) used, and the test bed or proof of concept employed by the authors. In this section, we present some key aspects of SFC implementations, where we focus on the open challenges (depicted in Figure 14) and highlight some of the existing efforts to solve the problems presented.

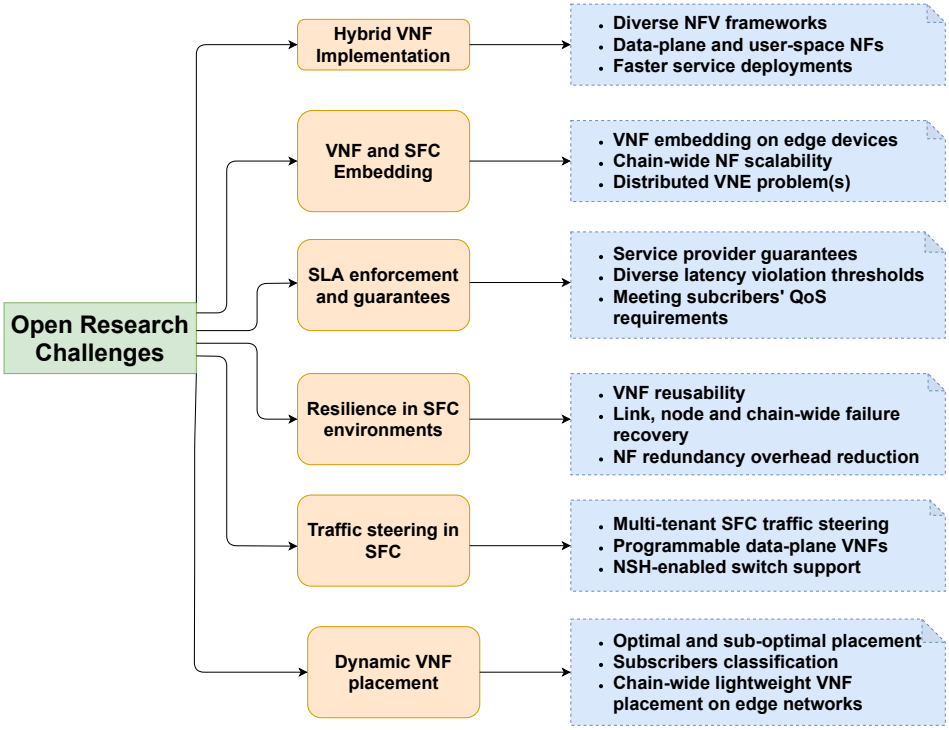

Figure 14. State-of-the-art and open research challenges.

5.1. Hybrid vNF Implementation

There are diverse frameworks for implementing virtual network functions in service provider network environments, which can be used to implement network functions at the data plane of the network, using virtual switching technologies such as OpenvSwitch [95] or P4 [53]. Virtual network functions can also be implemented in the user space, using unikernels [96], virtual machines [97], GPUs [98], or containers [13].

Because NFV and SDN allow for the creation of chains of network functions for efficient packet processing, it is ideal to create a processing pipeline that supports fast service delivery and efficient use of available resources by service providers to support future networks. Creating a service function chain (SFC) currently involves the chaining of network functions that are either implemented at the data plane or as user-space functions and are often carried out using network functions that are built using the same framework to create processing elements. The current data-plane technology does not easily provide support for operations such as the manipulation of floating-point values [99] or security modules.

Thus, the resource-constrained nature of the data plane also poses some limitations in terms of implementing resource-hungry network functions. We question the way in which network functions are currently implemented (including the composition of the network functions that create a service chain) and seek to argue for a hybrid implementation approach that combines user-space and data-plane components to create a chain of network functions along the processing pipeline.

It is correct to state that diverse frameworks have their constraints and benefits, and a hybrid implementation framework will leverage the benefits of both user space and data plane by composing processing pipelines using diverse frameworks. A hybrid NF implementation framework combines the available resources in the user space and the fast processing speed of the network data plane, with the aim of reducing end-to-end latency and improving packet throughput and CPU utilisation.

Van et al. [100] proposed a hybrid NFV framework for building low-latency and high-throughput vNFs using XDP. Simple operations are handled by the XDP program, whereas complex operations are handled by a user-space program. Three example NFs were chained (with two SFCs) using OpenStack. The performance of the framework was measured using throughput, latency, and CPU utilisation. They also considered the performance of the DPDK when interconnecting multiple vNFs. Their framework separates packet processing into fast and slow paths, with less complex NFs (simple LB and flow statistics) handled by the slow path, whereas vNFs that require complex processing are handled by the fast path. A use case example is a load balancer, which can be implemented in XDP if it uses a simple hash algorithm and is implemented using the hybrid approach if the NF is sophisticated.

Van et al. [101] used an extended Berkeley packet filter (eBPF), which is a Linux kernel framework that provides flexible kernel-level manipulation of packet processing pipelines. Their proposal was designed with the goal of having simple tasks handled by the kernel and complex NFs implemented in the user space. Two vNFs were built to test the proposed architecture: a dynamic traffic load balancer and a DPI network function. Compared to OpenFlow, the kernel-user interaction time is reduced with eBPF, which also provides better programmability options at the data plane.

Efforts such as HYPER [102] focus on creating a hybrid NFV framework that can leverage softwarized network functions and implementations on hardware devices, without considering whether the network functions are deployed at the network data plane or in the user space. Marcuzzo et al. [103] proposed a framework for offloading parts of a virtual network function to a programmable data plane.

The proposed architecture can offload specified components of a network function, and at the same time provide support in situations where NF offloading is not desirable. The management component of the framework comprises: (1) an interface for service providers (the user module), which can be used by users to initiate or stop an offload request; (2) a module for translating offload code that has been compiled for installation on data-plane programmable devices, (3) a module that serves as the offload manager,

for handling compiled offload code and communication with the NFV component; and (4) NFV and SDN modules for handling connections to the controller, topological data, and flow rule installation (SDN module), while communication with the offload agents on the network functions is handled by the NFV module.

Although the proposal presented by Marcuzzo et al. [103] aims to push some components of network functions down to the data plane of the network, it is not entirely a hybrid framework that composes packet processing pipelines from diverse NFV frameworks. Similarly, refs. [100–102] attempted to solve the problem of data plane and user-space packet processing using network functions built from the same framework(s).

5.2. vNF and SFC Embedding

Embedding virtual network functions on physical infrastructure is at the core of the creation of service function chains in service provider networks. Different authors have proposed solutions to solve the problems of vNF and SFC embedding. Reddy et al. [104] on embedding vNFs employed a mixed-integer linear programming (MILP) model, which handles the problem as a resource optimisation problem, where the scalability of the proposed model was enhanced using a variable neighbourhood search heuristic.

Pei et al. considered the embedding of SFCs in distributed cloud environments, with the aim of embedding service requests by reducing overheads. They approached the service chain embedding problem by formulating a binary integer programming (BIP) mechanism, which introduces algorithms aimed at the optimisation of the number of network functions that are placed along the service chain.

The survey of heuristic solutions for solving the virtual network embedding (VNE) problem presented by Cao et al. [105] captures some recent efforts that focus on embedding virtual network functions on substrate networks. Their work categorised heuristic algorithms for solving the VNE problem into (1) dynamic or static, (2) distributed or centralised, and (3) redundant or concise. Static VNE algorithms require less computation than dynamic solutions, which depends on the complexity of the substrate network and the virtual network function in use.

In their work, Sun et al. [106] proposed an energy-aware routing and adaptive delayed shutdown (EAR-ADS) model that supports the deployment of SFCs in a dynamic manner. Their proposed solution considers a practical scenario where overall deployment cost is minimized as well as balanced energy consumption by servers. Shutdown delays of servers are also reduced, which minimizes the effect of energy fluctuations associated with most network environments. The results presented show a huge reduction in energy costs and the overall stability of the network also improves.

The ability to dynamically reconfigure the embedded network functions makes the dynamic VNE algorithm ideal for next-generation networks [107]. Unlike the distributed VNE problem, the centralised problem makes use of a single-substrate infrastructure for embedding virtual network functions in future network environments. Computation in the distributed VNE algorithm is carried out by two or more substrate networks, which leads to improved scalability and the elimination of a single point of failure. Concise VNE heuristic algorithms are concerned with the exact number of substrate networks that are required to meet the SLA involved in embedding the virtual networks. The disadvantage of this approach is the lack of a guarantee for failure recovery. Redundant VNE solutions provide a provision for failure recovery by reserving the substrate network resources.

We believe that this remains an open challenge for the research community. Next-generation networks are generally envisaged to be completely softwarized, which should make the deployment of network functions doable at the network edge, which is closer to the point of traffic generation. The creation of chains of network functions is also seen as one of the major aspects of next-generation networks, which raises the question of embedding network functions on multiple commodity servers or embedding all functions on a single physical server for the creation of service chains. Questions such as the scalability of

such functions in the chain and the dynamic embedding of network functions, based on application profiles and QoS requirements, still need to be addressed.

5.3. SLA Enforcement and Guarantees

There is very little work currently in the literature that considers the enforcement of policies and SLAs by service providers, which revolves around finding out what service providers can guarantee in a NFV/SFC environment. Based on the recent literature, we argue that being able to adapt the QoS to frequent network changes is still a challenge in SFC environments for next-generation networks, and the ability to assess and visualise QoS requirements and parameters [108] will help in satisfying the long-term vision of SFC deployments.

Wang et al. [109] presented a QoE-driven service chain deployment, which also provides latency prediction features. Their focus was on improving the overall QoE for the user by reducing the number of rejections and waiting time experienced by users when accessing services. According to Herbaut et al. [110], Content Delivery Network (CDN) providers can negotiate the SLA with the vNF provider by requesting the creation of virtual CDN (vCDN) instances. An example is a CDN providing video streaming services such as YouTube or NetFlix. The focus of the SLA agreement between the CDN and the vNF provider in this scenario will be on the bandwidth and delay requirements of the service that is being provided.

Sun et al. [111] proposed SLA-NFV, which is an SLA-aware framework that focuses on the SLAs of the tenants (service subscribers). Their framework leverages a hybrid infrastructure, programmable hardware, and software, with the goal of enhancing the capability of NFV in handling different SLAs. The experimental results show that SLA-NFV, which creates a hybrid NFV, reduces latency by approximately 60% when compared to a software service chain. Their work did not consider a solution that helps service providers meet the diverse SLA requirements needed for the operation of next-generation networks.

5.4. Resilience in SFC Environments

Although our work considered efforts in the literature that try to solve the problem of resilience in NFV and SFC, resilience and fault tolerance are still major concerns in SFC environments, especially where the focus is on vNF reusability [112]. Achieving cost-effective resilience is still an open challenge in SFC implementation frameworks, and there are unanswered questions for future research which include how SFCs respond to failure conditions such as links between SFFs. The failure of the virtual network functions themselves is an aspect that needs to be considered as a long term vision that requires further research, that is, with the consideration of having redundant SFs along the SFP to handle failure scenarios.

Synthesising network processing graphs or service chains is another approach that is employed to achieve high availability in SFC, as well as the use of a multi-path routing approach [113]. As indicated by Mirjalily et al. [28], simple approaches such as the use of traffic load balancers can help with dynamic re-routing of traffic to alternate processing pipelines in SFC environments. However, an efficient load-balancing algorithm that tracks device states and available resources needs to be implemented, which goes beyond the basic round-robin algorithm that is commonly deployed in today's network environments.

Approaches such as the work by Ghaznavi et al. [88] try to provide SFC resilience and at the same time eliminate the need for NF replication, in order to reduce overhead. This is achieved by collecting and piggybacking NF state changes as packets traverse the service chain; thus, the overhead is reduced because the entire network function is not replicated to a standby node. A key challenge with proposals that employ the use of service replication is the amount of overhead incurred with redundant backup links, nodes, and service chains; thus, the efficient implementation of a high-availability failure mechanism is still a research challenge that requires further attention.

5.5. Traffic Steering in SFC

From the related literature on SFC traffic steering, we believe that the ability to dynamically steer traffic to the edge of the network, which adapts to network changes, is still required and desirable for the operation of next-generation networks. We also argue that having a functional programmable data-plane, which is a core component of future networks, can solve the problem of always having to install flow rules on virtual switches to steer traffic.

The ability to obtain the current network state and dynamically update the current service function path in real time requires more work to achieve better traffic steering in SFC for next-generation networks. Other factors such as delivery time, measuring virtual machines' (along the service path) usage, latency, and delay [114], are some of the considerations still open for further research with respect to service path selection in SFCs.

Our review of some related literature on SFC traffic steering also shows that the inter-operability between traffic steering techniques is still a challenge, as different service providers employ traffic steering approaches that best meet their business requirements. Multi-tenant networks also require a scalable traffic-steering scheme for SFC, as future networks are envisaged to be more heterogeneous in nature [115]. As it relates to steering traffic in SFC environments, service functions need to have support for SFC encapsulation protocols and headers.

Traffic steering types can be classified into three categories: header-based methods, tag-based methods, and programmable switch-based methods, which deal with the reclassification of flows and network isolation, where traffic forwarding depends on the configurations sent from the SDN controller to the switches. Tag-based steering methods make use of MAC addresses, VLANs, and MPLS tags to steer incoming traffic. Header-based methods use the network service header (NSH), service chain header (SCH), IP option field, and segmented routing header [27].

Medhat et al. [29] identified the absence of network service header (NSH) capability in switches as one of the challenges with traffic steering in SFC environments. Some proposals include the use of MAC addresses and tags to steer traffic in a service chain. Having virtual switches, such as Open vSwitch (OvS), which supports NSH capability, would allow for better traffic tagging and steering.

5.6. Dynamic vNF Placement

Most of the current approaches used in vNF placement, presented in Section 4, create a scenario in which all users contribute to a single latency violation threshold, which is common to all users. In line with the long term vision of next generation networks, an improved approach will help providers with different groups of applications such as VoIP/telephony users with a latency violation threshold, which should be lower than the threshold assigned to other classes of applications with less sensitive requirements in terms of latency. One justification that supports this argument is the much-anticipated rise in the number of end users with unique service requirements [57], which will affect the way in which vNFs are deployed in the future..

The placement of virtual Network Functions (vNFs) affects the latency between users and the vNFs, and thus we believe that a better placement design needs to be implemented for future networks, which prioritises the reduction in the negative effects of performance change, caused by the "hop-by-hop" movement of users between different vNFs. We argue that the cost implications of placing service functions in the SFC still need to be addressed, which is one of the open challenges in this domain. Other related challenges include creating placement schemes that consider parameters such as subscriber preferences, infrastructure properties, and delivery time.

Bhamare et al. [116] presented a novel fair weighted scheduling (FWS) solution for the scheduling of microservices in multi-cloud environments for the optimal creation of service function chains. Their proposed solution considers delays and SLA-related costs in deploying service function chains. They were able to reduce the overall turnaround time while

considering the total network delays and variable loads. The solution was compared to the standard biased greedy approach, which showed a notable increase in performance. The authors acknowledge the need for further investigations into microservices-related challenges such as security, fault tolerance, load balancing, and distributed data management.

Chai et al. [117] proposed a parallel placement scheme, PP-DRL, which uses deep reinforcement learning (DRL) to deploy SFCs optimally with minimal resource costs. They also used DRL to determine the right servers that could host service functions by collecting the characteristics of user requests as state information. They first used DRL to calculate the number of virtual machines and to find servers that can be used for hosting vNFs. The location of the end users is not prioritised in the scheme.

In their work, Laghrissi et al. [97] presented an efficient tool, which they called the "Network Slice Planner", for spatio-temporal simulation of mobile service usage, to maximise QoS for the users. They modified the classic predictive algorithm after presenting a set of existing placement algorithms. The performance of the enhanced predictive algorithm was compared with existing vNF placement algorithms, which showed slightly better results. Although they modelled the behaviour of the end users (in terms of traffic types), their proposed solution does not prioritise the placement of vNFs at the edge of the network.

Bhamare et al. [11] explain how the optimal placement of vNFs across multiple clouds is a problem which, when solved, can help in the optimization of parameters such as cost, network delay, and bandwidth. After presenting the components of an SFC environment, they presented an analytical model for the placement of service functions in multi-cloud network environments. Their work considers delays to end-users, QoS, and SLA. They employed an ILP approach to obtain the optimal solution, which was achieved by setting up an objective function and applicable constraints.

Cziva et al. [118] considered the placement of vNFs in a distributed-edge NFV environment with dynamic orchestration and re-calculation of vNF placement. They formulated and solved the edge vNF placement problem using the fundamentals of the optimal stopping theory (OST). They presented a time-optimized scheduler for optimal placement of vNFs at the edge of the network, and considered an undirected network graph comprising hosts, links between hosts, and users on the network. An assumption was made by the authors that resources on the hosts are finite and that links have a physical limit when it comes to bandwidth along the path.

6. Conclusions

The concept of service function chaining evolves as service providers continue to explore the benefits of deploying streamlined services for end users. Several approaches have been proposed in the literature for deploying network function virtualisation and solving different problems that involve chaining such network functions to meet user requirements.

In this work, we presented NFV environments and the requirements for chaining network functions for effective service delivery.

We conducted a comprehensive survey focusing on the NFV frameworks that also support the chaining of virtual network functions. We created a taxonomy of SFC implementation frameworks, a classification of the problems they attempted to solve, and discussed the open research challenges in SFC environments. The taxonomy presented separates the frameworks into three main categories: resource allocation and service orchestration, performance tuning, resilience, and fault recovery.

Important challenges that require further attention from the research community include the implementation of a hybrid NFV framework which can leverage processing resources from heterogeneous environments, including the orchestration of diverse vNFs. Other open challenges that should be addressed to support future networks are the optimisation of dynamic traffic steering in SFC, the efficient placement of heterogeneous network functions from diverse vendors, meeting SLA guarantees by service providers, and vNF/SFC embedding on commodity servers. All the open research challenges we

have presented describe the problem, the importance of addressing each problem, explain the subcomponents of the problem (also depicted in Figure 14), how each of the identified problems relate to the operation of next generation networks, and discuss prominent attempts to tackling them. Our work provides researchers in the NFV/SFC domain a clear picture of what has been achieved so far and the areas that require further research in order for the long term vision of service function chaining to be achieved.

Author Contributions: Conceptualization, H.U.A. and D.P. P.; Methodology, H.U.A. and D.P. P.; Resources, H.U.A. and D.P. P.; Writing—original draft preparation, H.U.A. and D.P. P.; Writing—review and editing, H.U.A. and D.P. P.; Supervision, D.P.P. All authors have read and agreed to this draft of the manuscript.

Funding: This work was supported in part by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/N033957/1, the PETRAS National Centre of Excellence for IoT Systems Cybersecurity, which has been funded by the UK EPSRC under grant number EP/S035362/1 and the Petroleum Technology Development Fund (PTDF) Nigeria, grant 1563/19.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: No conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

Abbreviation	Meaning
API	Application programming interface
CAPEX	Capital expenditure
CL	Classifier
CNF	Cloud-native function
DMA	Direct memory access
DPI	Deep packet inspection
eNF	embedded network function
ETSI	European Telecommunication Standards Institute
FPGA	Field-programmable gate array
EPC	Evolved packet core
ITU	International telecommunication union
IDS	Intrusion detection system
ILP	Integer linear programming
IoT	Internet of Things
IPS	Intrusion prevention system
PCAP	Packet capture
PoC	Proof of concept
RSS	Receive side scaling
SC	Service Chain
SDN	Software-defined networking
SFF	Service function forwarder
SP	Service provider

SF	Service function
SFCC	Service function chaining controller
MAC	Media access control
ML	Machine learning
NAT	Network address translation
NFV	Network function virtualization
NSH	Network service header
NFVI	NFV infrastructure
NFVO	NFV Orchestrator
NS	Network service
NFF	Network function forwarder
OF	OpenFlow
OPEX	Operational expenditure
QoS	Quality of service
QoE	Quality of experience
vCPE	virtual customer premises equipment
VNE	Virtual network embedding
VoIP	Voice over Internet Protocol
vNF	Virtual network function
VNFC	Virtual network function component
SFC	Service function chaining
VIM	Virtual infrastructure manager

References

1. Herrera, J.G.; Botero, J.F. Resource allocation in NFV: A comprehensive survey. *IEEE Trans. Netw. Serv. Manag.* **2016**, *13*, 518–532. [CrossRef]
2. Cherrared, S.; Imadali, S.; Fabre, E.; Gössler, G.; Yahia, I.G.B. A survey of fault management in network virtualization environments: Challenges and solutions. *IEEE Trans. Netw. Serv. Manag.* **2019**, *16*, 1537–1551. [CrossRef]
3. Paganelli, F.; Cappanera, P.; Cuffaro, G. Tenant-defined service function chaining in a multi-site network slice. *Future Gener. Comput. Syst.* **2021**, *121*, 1–18. [CrossRef]
4. Laghrissi, A.; Taleb, T. A survey on the placement of virtual resources and virtual network functions. *IEEE Commun. Surv. Tutor.* **2018**, *21*, 1409–1434. [CrossRef]
5. Bujari, A.; Palazzi, C.E.; Polonio, D.; Zanella, M. Service function chaining: A lightweight container-based management and orchestration plane. In Proceedings of the 2019 16th IEEE Annual Consumer Communications & Networking Conference (CCNC), Las Vegas, NV, USA, 11–14 January 2019; pp. 1–4.
6. Santos, J.; Wauters, T.; Volckaert, B.; De Turck, F. Towards delay-aware container-based service function chaining in fog computing. In Proceedings of the NOMS 2020-2020 IEEE/IFIP Network Operations and Management Symposium, Budapest, Hungary, 20–24 April 2020; pp. 1–9.
7. Li, X.; Qian, C. A survey of network function placement. In Proceedings of the 2016 13th IEEE Annual Consumer Communications & Networking Conference (CCNC), Las Vegas, NV, USA, 9–12 January 2016; pp. 948–953.
8. Haleplidis, E.; Joachimpillai, D.; Salim, J.H.; Lopez, D.; Martin, J.; Pentikousis, K.; Denazis, S.; Koufopavlou, O. ForCES applicability to SDN-enhanced NFV. In Proceedings of the 2014 Third European Workshop on Software Defined Networks, Budapest, Hungary, 1–3 September 2014; pp. 43–48.
9. Liu, J.; Lu, W.; Zhou, F.; Lu, P.; Zhu, Z. On dynamic service function chain deployment and readjustment. *IEEE Trans. Netw. Serv. Manag.* **2017**, *14*, 543–553. [CrossRef]
10. Chua, F.C.; Ward, J.; Zhang, Y.; Sharma, P.; Huberman, B.A. Stringer: Balancing latency and resource usage in service function chain provisioning. *IEEE Internet Comput.* **2016**, *20*, 22–31. [CrossRef]
11. Bhamare, D.; Samaka, M.; Erbad, A.; Jain, R.; Gupta, L.; Chan, H.A. Optimal virtual network function placement in multi-cloud service function chaining architecture. *Comput. Commun.* **2017**, *102*, 1–16. [CrossRef]
12. Martins, J.; Ahmed, M.; Raiciu, C.; Olteanu, V.; Honda, M.; Bifulco, R.; Huici, F. ClickOS and the art of network function virtualization. In Proceedings of the 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14), Seattle, WA, USA, 2–4 April 2014; pp. 459–473.
13. Zhang, W.; Liu, G.; Zhang, W.; Shah, N.; Lopreiato, P.; Todeschi, G.; Ramakrishnan, K.; Wood, T. OpenNetVM: A platform for high performance network service chains. In Proceedings of the 2016 Workshop on Hot topics in Middleboxes and Network Function Virtualization, Florianopolis, Brazil, 22 August 2016; pp. 26–31.
14. Bremler-Barr, A.; Harchol, Y.; Hay, D. OpenBox: A software-defined framework for developing, deploying, and managing network functions. In Proceedings of the 2016 ACM SIGCOMM Conference, Florianopolis, Brazil, 22 August 2016; pp. 511–524.
15. Anwer, B.; Benson, T.; Feamster, N.; Levin, D. Programming slick network functions. In Proceedings of the 1st ACM Sigcomm Symposium on Software Defined Networking Research, Santa Clara, CA, USA, 17 June 2015; pp. 1–13.

16. Katsikas, G.P.; Enguehard, M.; Kuźniar, M.; Maguire Jr, G.Q.; Kostić, D. SNF: Synthesizing high performance NFV service chains. *PeerJ Comput. Sci.* **2016**, *2*, e98. [CrossRef]
17. Katsikas, G.P.; Barbette, T.; Kostic, D.; Steinert, R.; Maguire, G.Q., Jr. Metron: NFV Service Chains at the True Speed of the Underlying Hardware. In Proceedings of the 15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18), Renton, WA, USA, 9–11 April 2018; pp. 171–186.
18. Panda, A.; Han, S.; Jang, K.; Walls, M.; Ratnasamy, S.; Shenker, S. NetBricks: Taking the V out of NFV. In Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), Savannah, GA, USA, 2 November 2016; pp. 203–216.
19. Cziva, R.; Jouet, S.; White, K.J.; Pezaros, D.P. Container-based network function virtualization for software-defined networks. In Proceedings of the 2015 IEEE Symposium on Computers and Communication (ISCC), Larnaca, Cyprus, 6–9 July 2015; pp. 415–420.
20. Zoure, M.; Ahmed, T.; Réveillére, L. Network Services Anomalies in NFV: Survey, Taxonomy, and Verification Methods. *IEEE Trans. Netw. Serv. Manag.* **2022**, doi:10.1109/TNSM.2022.3144582. [CrossRef]
21. Zhang, T.; Qiu, H.; Linguaglossa, L.; Cerroni, W.; Giaccone, P. NFV platforms: Taxonomy, design choices and future challenges. *IEEE Trans. Netw. Serv. Manag.* **2020**, *18*, 30–48. [CrossRef]
22. Hamdan, M.; Hassan, E.; Abdelaziz, A.; Elhigazi, A.; Mohammed, B.; Khan, S.; Vasilakos, A.V.; Marsono, M.N. A comprehensive survey of load balancing techniques in software-defined network. *J. Netw. Comput. Appl.* **2021**, *174*, 102856. [CrossRef]
23. Fei, X.; Liu, F.; Zhang, Q.; Jin, H.; Hu, H. Paving the Way for NFV Acceleration: A Taxonomy, Survey and Future Directions. *ACM Comput. Surv. (CSUR)* **2020**, *53*, 1–42. [CrossRef]
24. Kaur, K.; Mangat, V.; Kumar, K. A comprehensive survey of service function chain provisioning approaches in SDN and NFV architecture. *Comput. Sci. Rev.* **2020**, *38*, 100298. [CrossRef]
25. Hantouti, H.; Benamar, N.; Taleb, T. Service Function Chaining in 5G and Beyond Networks: Challenges and Open Research Issues. *IEEE Netw.* **2020**, *34*, 320–327. [CrossRef]
26. Bonfim, M.S.; Dias, K.L.; Fernandes, S.F. Integrated NFV/SDN architectures: A systematic literature review. *ACM Comput. Surv. (CSUR)* **2019**, *51*, 1–39. [CrossRef]
27. Hantouti, H.; Benamar, N.; Taleb, T.; Laghrissi, A. Traffic steering for service function chaining. *IEEE Commun. Surv. Tutor.* **2018**, *21*, 487–507. [CrossRef]
28. Mirjalily, G.; Zhiquan, L. Optimal network function virtualization and service function chaining: A survey. *Chin. J. Electron.* **2018**, *27*, 704–717. [CrossRef]
29. Medhat, A.M.; Taleb, T.; Elmangoush, A.; Carella, G.A.; Covaci, S.; Magedanz, T. Service function chaining in next generation networks: State of the art and research challenges. *IEEE Commun. Mag.* **2016**, *55*, 216–223. [CrossRef]
30. Bera, S.; Misra, S.; Vasilakos, A.V. Software-defined networking for internet of things: A survey. *IEEE Internet Things J.* **2017**, *4*, 1994–2008. [CrossRef]
31. Veeraraghavan, M.; Sato, T.; Buchanan, M.; Rahimi, R.; Okamoto, S.; Yamanaka, N. Network function virtualization: A survey. *IEICE Trans. Commun.* **2017**, E100B, 1978–1991. [CrossRef]
32. Bhamare, D.; Jain, R.; Samaka, M.; Erbad, A. A survey on service function chaining. *J. Netw. Comput. Appl.* **2016**, *75*, 138–155. [CrossRef]
33. Xie, Y.; Liu, Z.; Wang, S.; Wang, Y. Service function chaining resource allocation: A survey. *arXiv* **2016**, arXiv:1608.00095.
34. Yang, M.; Li, Y.; Jin, D.; Zeng, L.; Wu, X.; Vasilakos, A.V. Software-defined and virtualized future mobile and wireless networks: A survey. *Mob. Netw. Appl.* **2015**, *20*, 4–18. [CrossRef]
35. Quinn, P.; Beliveau, A. Service Function Chaining (SFC) Architecture. draft-quinn-sfc-arch-04. 2014. Available online: https://www.ietf.org/proceedings/89/slides/slides-89-sfc-10.pdf (accessed on 20 January 2022).
36. Gasparakis, J.; Smith, K.; Zhou, D. Evaluating Dynamic Service Function Chaining for the Gi-LAN. In *White Paper*; Intel: Santa Clara, CA, USA, 2016.
37. Halpern, J.; Pignataro, C. Service Function Chaining (sfc) Architecture. In *RFC 7665*; IETF: 2015; pp. 1 –28. Available online: https://www.hjp.at/(de)/doc/rfc/rfc7665.html (accessed on 20 January 2022).
38. Grønsund, P.; Mahmood, K.; Millstein, G.; Noy, A.; Solomon, G.; Sahai, A. A solution for SGi-LAN services virtualization using NFV and SDN. In Proceedings of the 2015 European Conference on Networks and Communications (EuCNC), Paris, France, 29 June–2 July 2015; pp. 408–412.
39. Naik, P.; Vutukuru, M. libVNF: A Framework for Building Scalable High Performance Virtual Network Functions. In Proceedings of the 8th Asia-Pacific Workshop on Systems, Mumbai, India, 2 September 2017; pp. 1–8.
40. Turk, Y.; Zeydan, E. An Implementation of Network Service Chaining for SDN-enabled Mobile Packet Data Networks. In Proceedings of the 2019 International Symposium on Networks, Computers and Communications (ISNCC), Istanbul, Turkey, 18–20 June 2019; pp. 1–6.
41. KAUR, K.; KUMAR, K.; MANGAT, V. A road to network function virtualization and applications. *Adv. Math. Sci. J.* **2020**, *9*, 4059–4066. [CrossRef]
42. Brown, G.; Reading, H. Service Chaining in Carrier Networks. *Heavy Read.* **2015**. Available online: https://www.qosmos.com/wp-content/uploads/Service-Chaining-in-Carrier-Networks_WP_Heavy-Reading_Qosmos_Feb2015.pdf (accessed on 20 January 2022).

43. Shojafar, M.; Pooranian, Z.; Sookhak, M.; Buyya, R. Recent advances in cloud data centers toward fog data centers. *Concurr. Comput. Pract. Exp.* **2019**, *31*, e5164. [CrossRef]
44. Cunha, V.A.; Cardoso, I.D.; Barraca, J.P.; Aguiar, R.L. Policy-driven vCPE through dynamic network service function chaining. In Proceedings of the 2016 IEEE NetSoft Conference and Workshops (NetSoft), Seoul, Korea (South), 6–10 June 2016; pp. 156–160.
45. Yan, Z.; Zhang, P.; Vasilakos, A.V. A security and trust framework for virtualized networks and software-defined networking. *Secur. Commun. Netw.* **2016**, *9*, 3059–3069. [CrossRef]
46. Liu, Y.; Zhou, F.; Chen, C.; Zhu, Z.; Shang, T.; Torres-Moreno, J.M. Disaster protection in Inter-DataCenter networks leveraging cooperative storage. *IEEE Trans. Netw. Serv. Manag.* **2021**, *18*, 2598–2611. [CrossRef]
47. Zhong, X.; Wang, Y.; Qiu, X. Service function chain orchestration across multiple clouds. *China Commun.* **2018**, *15*, 99–116. [CrossRef]
48. Medhat, A.M.; Carella, G.A.; Pauls, M.; Monachesi, M.; Corici, M.; Magedanz, T. Resilient orchestration of Service Functions Chains in a NFV environment. In Proceedings of the 2016 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), Palo Alto, CA, USA , 7–10 November 2016; pp. 7–12.
49. Sarmiento, D.E.; Lebre, A.; Nussbaum, L.; Chari, A. Decentralized SDN Control Plane for a Distributed Cloud-Edge Infrastructure: A Survey. *IEEE Commun. Surv. Tutor.* **2021**, *23*, 256–281. [CrossRef]
50. Medved, J.; Varga, R.; Tkacik, A.; Gray, K. Opendaylight: Towards a model-driven sdn controller architecture. In Proceeding of the IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks 2014, Sydney, NSW, Australia, 19 June 2014; pp. 1–6.
51. Kaur, S.; Singh, J.; Ghumman, N.S. Network programmability using POX controller. In Proceeding of the International Conference on Communication, Computing & Systems (ICCCS), 2014; Volume 138, pp. 134–138. Available online: https://docplayer.net/11300937-Network-programmability-using-pox-controller.html (accessed on 20 January 2022).
52. Zhang, T.; Linguaglossa, L.; Giaccone, P.; Iannone, L.; Roberts, J. Performance benchmarking of state-of-the-art software switches for NFV. *Comput. Netw.* **2021**, *188*, 107861. [CrossRef]
53. Bosshart, P.; Daly, D.; Gibb, G.; Izzard, M.; McKeown, N.; Rexford, J.; Schlesinger, C.; Talayco, D.; Vahdat, A.; Varghese, G.; et al. P4: Programming protocol-independent packet processors. *ACM SIGCOMM Comput. Commun. Rev.* **2014**, *44*, 87–95. [CrossRef]
54. ETSI. *Network Functions Virtualisation (NFV): Architectural Framework*; Technical Report 002 V1.1.1; 2013. Available online: https://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.01.01_60/gs_NFV002v010101p.pdf (accessed on 20 January 2022).
55. Binu, A.; Kumar, G.S. Virtualization techniques: A methodical review of XEN and KVM. In *International Conference on Advances in Computing and Communications*; Springer: Berlin/Heidelberg, Germany, 2011; pp. 399–410.
56. Wray, M.J.; Dalton, C.I. Network Virtualization. US Patent 8,223,770, 17 July 2012. Available online: https://uspto.report/patent/grant/8,223,770 (accessed on 20 January 2022).
57. Wang, M.; Cheng, B.; Wang, S.; Chen, J. Availability-and traffic-aware placement of parallelized SFC in data center networks. *IEEE Trans. Netw. Serv. Manag.* **2021**, *18*, 182–194. [CrossRef]
58. Özdem, M.; Alkan, M. Subscriber aware dynamic service function chaining. *Comput. Netw.* **2021**, *194*, 108138. [CrossRef]
59. Li, Y.; Chen, M. Software-defined network function virtualization: A survey. *IEEE Access* **2015**, *3*, 2542–2553.
60. Yousaf, F.Z.; Bredel, M.; Schaller, S.; Schneider, F. NFV and SDN—Key technology enablers for 5G networks. *IEEE J. Sel. Areas Commun.* **2017**, *35*, 2468–2478. [CrossRef]
61. Kak, A. Towards 6G Through SDN and NFV-Based Solutions for Terrestrial and Non-Terrestrial Networks. Ph.D. Thesis, Georgia Institute of Technology, Atlanta, GA, USA, 2021.
62. Qadri, Y.A.; Nauman, A.; Zikria, Y.B.; Vasilakos, A.V.; Kim, S.W. The future of healthcare internet of things: A survey of emerging technologies. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 1121–1167. [CrossRef]
63. Huang, M.; Liu, A.; Xiong, N.N.; Wang, T.; Vasilakos, A.V. An effective service-oriented networking management architecture for 5G-enabled internet of things. *Comput. Netw.* **2020**, *173*, 107208. [CrossRef]
64. Morocho-Cayamcela, M.E.; Lee, H.; Lim, W. Machine learning for 5G/B5G mobile and wireless communications: Potential, limitations, and future directions. *IEEE Access* **2019**, *7*, 137184–137206. [CrossRef]
65. Berardinelli, G.; Mahmood, N.H.; Rodriguez, I.; Mogensen, P. Beyond 5G wireless IRT for industry 4.0: Design principles and spectrum aspects. In Proceedings of the 2018 IEEE Globecom Workshops (GC Wkshps), Abu Dhabi, United Arab Emirates, 9–13 December 2018; pp. 1–6.
66. Katz, M.; Matinmikko-Blue, M.; Latva-Aho, M. 6Genesis flagship program: Building the bridges towards 6G-enabled wireless smart society and ecosystem. In Proceedings of the 2018 IEEE 10th Latin-American Conference on Communications (LATINCOM), Guadalajara, Mexico, 14–16 November 2018; pp. 1–9.
67. Abdelwahab, S.; Hamdaoui, B.; Guizani, M.; Znati, T. Network function virtualization in 5G. *IEEE Commun. Mag.* **2016**, *54*, 84–91. [CrossRef]
68. Huang, H.; Zeng, C.; Zhao, Y.; Min, G.; Zhu, Y.Y.; Miao, W.; Hu, J. Scalable Service Function Chain Orchestration in NFV-enabled Networks: A Federated Reinforcement Learning Approach. *IEEE J. Sel. Areas Commun.* **2021**, *39*, 2558–2571. [CrossRef]
69. Sun, G.; Xu, Z.; Yu, H.; Chen, X.; Chang, V.; Vasilakos, A.V. Low-latency and resource-efficient service function chaining orchestration in network function virtualization. *IEEE Internet Things J.* **2019**, *7*, 5760–5772. [CrossRef]
70. Ballani, H.; Costa, P.; Gkantsidis, C.; Grosvenor, M.P.; Karagiannis, T.; Koromilas, L.; O'Shea, G. Enabling end-host network functions. *ACM SIGCOMM Comput. Commun. Rev.* **2015**, *45*, 493–507. [CrossRef]

71. Palkar, S.; Lan, C.; Han, S.; Jang, K.; Panda, A.; Ratnasamy, S.; Rizzo, L.; Shenker, S. E2: A framework for NFV applications. In Proceedings of the 25th Symposium on Operating Systems Principles, Monterey, CA, USA, 4 October 2015; pp. 121–136.
72. Carella, G.A.; Magedanz, T. Open baton: A framework for virtual network function management and orchestration for emerging software-based 5G networks. *Newsletter* **2015**, *2016*, 190.
73. Kouchaksaraei, H.R.; Dierich, T.; Karl, H. Pishahang: Joint orchestration of network function chains and distributed cloud applications. In Proceedings of the 2018 4th IEEE Conference on Network Softwarization and Workshops (NetSoft), Montreal, QC, Canada, 25–29 June 2018; pp. 344–346.
74. Mafioletti, D.R.; Dominicini, C.K.; Martinello, M.; Ribeiro, R.M.; Villaca, R.d.S. PIaFFE: A Place-as-you-go In-network Framework for Flexible Embedding of VNFs. In Proceedings of the IEEE International Conference on Communications, Dublin, Ireland, 7–11 June 2020.
75. Kouchaksaraei, H.R.; Karl, H. Service Function Chaining Across OpenStack and Kubernetes Domains. In Proceedings of the 13th ACM International Conference on Distributed and Event-Based Systems, Darmstadt, Germany, 24 June 2019; pp. 240–243.
76. Dab, B.; Fajjari, I.; Rohon, M.; Auboin, C.; Diquélou, A. An Efficient Traffic Steering for Cloud-Native Service Function Chaining. In Proceedings of the 2020 23rd Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN), Paris, France, 24–27 February 2020; pp. 71–78.
77. Kohler, E.; Morris, R.; Chen, B.; Jannotti, J.; Kaashoek, M.F. The Click modular router. *ACM Trans. Comput. Syst. (TOCS)* **2000**, *18*, 263–297. [CrossRef]
78. Katsikas, G.P.; Barbette, T.; Kostić, D.; Maguire, J.G.Q.; Steinert, R. Metron: High-performance NFV Service Chaining Even in the Presence of Blackboxes. *ACM Trans. Comput. Syst. (TOCS)* **2021**, *38*, 1–45. [CrossRef]
79. Meng, Z.; Bi, J.; Wang, H.; Sun, C.; Hu, H. CoCo: Compact and optimized consolidation of modularized service function chains in NFV. In Proceedings of the 2018 IEEE International Conference on Communications (ICC), Kansas City, MO, USA, 20–24 May 2018; pp. 1–7.
80. Li, L.; Ota, K.; Dong, M. DeepNFV: A lightweight framework for intelligent edge network functions virtualization. *IEEE Netw.* **2018**, *33*, 136–141. [CrossRef]
81. Lombardo, A.; Manzalini, A.; Schembra, G.; Faraci, G.; Rametta, C.; Riccobene, V. An open framework to enable NetFATE (Network Functions at the edge). In Proceedings of the 2015 1st IEEE Conference on Network Softwarization (NetSoft), London, UK, 13–17 April 2015; pp. 1–6.
82. Yasukata, K.; Huici, F.; Maffione, V.; Lettieri, G.; Honda, M. HyperNF: Building a high performance, high utilization and fair NFV platform. In Proceedings of the 2017 Symposium on Cloud Computing, Santa Clara, CA, USA, 24 September 2017; pp. 157–169.
83. Castanho, M.S.; Dominicini, C.K.; Villacça, R.S.; Martinello, M.; Ribeiro, R.M. Phantomsfc: A fully virtualized and agnostic service function chaining architecture. In Proceedings of the 2018 IEEE Symposium on Computers and Communications (ISCC), Natal, Brazil, 25–28 June 2018; pp. 354–359.
84. Zheng, C.; Lu, Q.; Li, J.; Liu, Q.; Fang, B. A flexible and efficient container-based nfv platform for middlebox networking. In Proceedings of the 33rd Annual ACM Symposium on Applied Computing, Pau, France, 9 April 2018; pp. 989–995.
85. Meng, Z.; Bi, J.; Wang, H.; Sun, C.; Hu, H. MicroNF: An efficient framework for enabling modularized service chains in NFV. *IEEE J. Sel. Areas Commun.* **2019**, *37*, 1851–1865. [CrossRef]
86. Eisenbud, D.E.; Yi, C.; Contavalli, C.; Smith, C.; Kononov, R.; Mann-Hielscher, E.; Cilingiroglu, A.; Cheyney, B.; Shang, W.; Hosein, J.D. Maglev: A fast and reliable software network load balancer. In Proceedings of the 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16), Santa Clara, CA, USA, 16–18 March 2016; pp. 523–535.
87. Quinn, P.; Elzur, U.; Pignataro, C. Network Service Header (NSH). In *RFC 8300*; 2018; pp. 1–40. Available online: https://www.hjp.at/doc/rfc/rfc8300.html (accessed on 20 January 2022).
88. Ghaznavi, M.; Jalalpour, E.; Wong, B.; Boutaba, R.; Mashtizadeh, A.J. Fault tolerant service function chaining. In Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication, Virtual Event, USA, 30 July 2020; pp. 198–210.
89. Wang, L.; Mao, W.; Zhao, J.; Xu, Y. DDQP: A double deep Q-learning approach to online fault-tolerant SFC placement. *IEEE Trans. Netw. Serv. Manag.* **2021**, *18*, 118–132. [CrossRef]
90. Kulkarni, S.G.; Liu, G.; Ramakrishnan, K.; Arumaithurai, M.; Wood, T.; Fu, X. REINFORCE: Achieving Efficient Failure Resiliency for Network Function Virtualization-Based Services. *IEEE/ACM Trans. Netw.* **2020**, *28*, 695–708. [CrossRef]
91. Hmaity, A.; Savi, M.; Musumeci, F.; Tornatore, M.; Pattavina, A. Virtual network function placement for resilient service chain provisioning. In Proceedings of the 2016 8th International Workshop on Resilient Networks Design and Modeling (RNDM), Halmstad, Sweden, 13–15 September 2016; pp. 245–252.
92. Harchol, Y.; Hay, D.; Orenstein, T. FTvNF: Fault tolerant virtual network functions. In Proceedings of the 2018 Symposium on Architectures for Networking and Communications Systems, Ithaca, NY, USA, 23 July 2018; pp. 141–147.
93. Nguyen, H.B.; Dinh, N.T.; Oh, J.; Kim, Y. An Openflow-based Scheme for Service Chaining's High Availability in Cloud Network. In Proceedings of the International Conference on ICT Convergence, Jeju, Korea (South), 16–18 October 2019.
94. Sherry, J.; Gao, P.X.; Basu, S.; Panda, A.; Krishnamurthy, A.; Maciocco, C.; Manesh, M.; Martins, J.; Ratnasamy, S.; Rizzo, L.; et al. Rollback-recovery for middleboxes. In Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, London United Kingdom, 17 August 2015; pp. 227–240.

95. Jackson, E.J.; Walls, M.; Panda, A.; Pettit, J.; Pfaff, B.; Rajahalme, J.; Koponen, T.; Shenker, S. SoftFlow: A Middlebox Architecture for Open vSwitch. In Proceedings of the USENIX Annual Technical Conference (ATC 16), Berkeley, CA, USA, 22–24 June 2016; pp. 15–28. Available online: https://www.usenix.org/conference/atc16/technical-sessions/presentation/jackson (accessed on 20 January 2022).
96. Kurek, T. Unikernel Network Functions: A Journey Beyond the Containers. *IEEE Commun. Mag.* **2019**, *57*, 15–19. [CrossRef]
97. Laghrissi, A.; Taleb, T.; Bagaa, M.; Flinck, H. Towards edge slicing: VNF placement algorithms for a dynamic & realistic edge cloud environment. In Proceedings of the IEEE Global Communications Conference (GLOBECOM), Singapore, Singapore, 4–8 December 2017; pp. 1–6.
98. Yi, X.; Duan, J.; Wu, C. Gpunfv: A gpu-accelerated nfv system. In Proceedings of the First Asia-Pacific Workshop on Networking, Hong Kong, China, 3 August 2017; pp. 85–91.
99. Simpson, K.A.; Cziva, R.; Pezaros, D.P. Seiðr: Dataplane Assisted Flow Classification Using ML. 2020. Available online: https://ieeexplore.ieee.org/abstract/document/9348063 (accessed on 20 January 2022).
100. Van Tu, N.; Yoo, J.H.; Hong, J.W.K. Building hybrid virtual network functions with eXpress data path. In Proceedings of the 2019 15th International Conference on Network and Service Management (CNSM), Halifax, NS, Canada, 21–25 October 2019; pp. 1–9.
101. Van Tu, N.; Ko, K.; Hong, J.W.K. Architecture for building hybrid kernel-user space virtual network functions. In Proceedings of the 2017 13th International Conference on Network and Service Management (CNSM), Tokyo, Japan, 26–30 November 2017; pp. 1–6.
102. Sun, C.; Bi, J.; Zheng, Z.; Hu, H. HYPER: A hybrid high-performance framework for network function virtualization. *IEEE J. Sel. Areas Commun.* **2017**, *35*, 2490–2500. [CrossRef]
103. Marcuzzo, L.d.C.; dos Santos, C.R. Enabling Partial Offload of Virtualized Network Functions into the Programmable Data Plane. In Proceedings of the 2020 IEEE Latin-American Conference on Communications (LATINCOM), Santo Domingo, Dominican Republic, 18–20 November 2020; pp. 1–6.
104. Reddy, V.S.; Baumgartner, A.; Bauschert, T. Robust embedding of VNF/service chains with delay bounds. In Proceedings of the 2016 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), Palo Alto, CA, USA, 7–9 November 2016; pp. 93–99.
105. Cao, H.; Hu, H.; Qu, Z.; Yang, L. Heuristic solutions of virtual network embedding: A survey. *China Commun.* **2018**, *15*, 186–219. [CrossRef]
106. Sun, G.; Zhou, R.; Sun, J.; Yu, H.; Vasilakos, A.V. Energy-efficient provisioning for service function chains to support delay-sensitive applications in network function virtualization. *IEEE Internet Things J.* **2020**, *7*, 6116–6131. [CrossRef]
107. Li, J.; Shi, W.; Ye, Q.; Zhang, N.; Zhuang, W.; Shen, X. Multi-service function chain embedding with delay-guarantee: A game-theoretical approach. *IEEE Internet Things J.* **2021**, *8*, 11219–11232. [CrossRef]
108. Zhao, D.; Luo, L.; Yu, H.; Chang, V.; Buyya, R.; Sun, G. Security-SLA-guaranteed service function chain deployment in cloud-fog computing networks. *Cluster Comput.* **2021**, *24*, 2479–2494. [CrossRef]
109. Wang, I.C.; Wen, C.H.P.; Chao, H.J. Improving Quality of Experience of Service-Chain Deployment for Multiple Users. In Proceedings of the 2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS), Banff, AB, Canada, 4–6 June 2018; pp. 1–6.
110. Herbaut, N.; Negru, D.; Magoni, D.; Frangoudis, P.A. Deploying a content delivery service function chain on an SDN-NFV operator infrastructure. In Proceedings of the 2016 International Conference on Telecommunications and Multimedia (TEMU), Heraklion, Greece, 25–27 July 2016; pp. 1–7.
111. Sun, C.; Bi, J.; Zheng, Z.; Hu, H. Sla-nfv: An sla-aware high performance framework for network function virtualization. In Proceedings of the 2016 ACM SIGCOMM Conference, Florianopolis, Brazil, 22 August 2016; pp. 581–582.
112. Chowdhury, S.R.; Salahuddin, M.A.; Limam, N.; Boutaba, R. Re-architecting NFV ecosystem with microservices: State of the art and research challenges. *IEEE Netw.* **2019**, *33*, 168–176. [CrossRef]
113. Cai, S.; Zhou, F.; Zhang, Z.; Meddahi, A. Disaster-Resilient Service Function Chain Embedding Based on Multi-Path Routing. In Proceedings of the IEEE INFOCOM 2021-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), Vancouver, BC, Canada, 10–13 May 2021; pp. 1–7.
114. Németh, B.; Molner, N.; Martinperez, J.; Bernardos, C.J.; De la Oliva, A.; Sonkoly, B. Delay and reliability-constrained VNF placement on mobile and volatile 5G infrastructure. *IEEE Trans. Mob. Comput.* **2021**, doi:10.1109/TMC.2021.3055426. [CrossRef]
115. Bouridah, A.; Fajjari, I.; Aitsaadi, N.; Belhadef, H. Optimized Scalable SFC Traffic Steering Scheme for Cloud Native based Applications. In Proceedings of the 2021 IEEE 18th Annual Consumer Communications & Networking Conference (CCNC), Las Vegas, NV, USA, 9–12 January 2021; pp. 1–6.
116. Bhamare, D.; Samaka, M.; Erbad, A.; Jain, R.; Gupta, L. Exploring microservices for enhancing internet QoS. *Trans. Emerg. Telecommun. Technol.* **2018**, *29*, e3445. [CrossRef]
117. Chai, H.; Zhang, J.; Wang, Z.; Shi, J.; Huang, T. A Parallel Placement Approach for Service Function Chain Using Deep Reinforcement Learning. In Proceedings of the 2019 IEEE 5th International Conference on Computer and Communications (ICCC), Chengdu, China, 6–9 December 2019; pp. 2123–2128.
118. Cziva, R.; Anagnostopoulos, C.; Pezaros, D.P. Dynamic, latency-optimal VNF placement at the network edge. In Proceedings of the IEEE Infocom 2018-IEEE Conference on Computer Communications, Honolulu, HI, USA, 16–19 April 2018; pp. 693–701.