Computer Networks 241 (2024) 110202

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

5G-VIOS: Towards next generation intelligent inter-domain network service orchestration and resource optimisation

Shadi Moazzeni [a,*], Konstantinos Katsaros [b], Nasim Ferdosian [c], Konstantinos Antonakoglou [b], Mark Rouse [b], Dritan Kaleshi [b], Adriana Fernández-Fernández [d], Miguel Catalan-Cid [d], Constantinos Vrontos [a], Reza Nejabati [a], Dimitra Simeonidou [a,1]

[a] Smart Internet Lab, University of Bristol, UK
[b] Digital Catapult, UK
[c] Cisco-Curtin Centre for Networks, Curtin University, Bentley, Australia
[d] i2CAT Foundation, Spain

ARTICLE INFO

Keywords:
Cross-domain orchestration
VNF profiling
Resource minimising

ABSTRACT

This paper introduces an intelligent network service orchestration platform, referred to as 5G-VIOS for 5G networks and beyond in accordance with the Zero-touch Network and Service Management (ZSM) paradigm. The proposed solution is responsible for an automated network slicing and life-cycle management of network applications and services, across multiple administrative and technological domains. An Artificial Intelligent (AI) model utilising Machine Learning (ML) techniques is exploited to intelligently and efficiently profile the network services and predict the efficient configuration of resources needed to meet the performance targets and Service Level Agreements (SLAs) of these network services across multiple domains. We test and validate the performance of the prediction models for both resource configuration and utilisation in various settings for different resources and data rates. We also showcase how resource utilisation predictions of a virtualised network service can significantly assist in its life cycle management by proactively preventing unnecessary actions such as its migration.

1. Introduction

The advances in mobile network technologies with the advent of 5G have introduced new capabilities supporting numerous diverse use-cases requiring high throughput, ultra-low latency, and high connection density, which are not achieved by the current one-size-fits-all network designs. This necessitates upgrading the network with the capability of flexible autonomous service deployment and on-demand networking for different use-cases.

Network slicing is considered a promising solution for creating service-customised 5G networks that efficiently harness the capabilities of evolving technologies such as Software-Defined Networking (SDN) and Network Function Virtualisation (NFV). Network slicing exploits programmability and modularity during allocating the network resources to specific vertical service requirements and thus, transforms network resource provisioning from one-size-fits-all to one-size-per-service methodology. Each network slice instance works as an end-to-end logical "dedicated" network over the same underlying physical networks. Network slices, customised and allocated to a specific use-case, can stretch across larger geographical areas with multiple administrative domains.

Deploying a slice across multiple domains is challenging not only from the perspective of decomposing the slice request into the respective administrative domain but also for guaranteeing its performance [1]. This necessitates designing a MANagement and Orchestration (MANO) platform that can deploy inter-domain network slices. This paper mainly addresses the scenarios being investigated by the Horizon2020 5G-VICTORI Project [2], specifically leveraging the infrastructure and resources available at different facilities for delivering *Network Slice as a Service (NSaaS)* on demand to the user. The functional architecture of the project is illustrated in Fig. 1. The 5G-VICTORI platform develops an open data management platform for scalable data collection, aggregation and processing across the various project infrastructure sites, adopting ML and Artificial Intelligence (AI) techniques to offer optimised vertical services.

A thin inter-domain orchestration layer on top of the orchestration solutions of the individual facility sites is developed to enable dynamic

* Corresponding author.
E-mail address: shadi.moazzeni@bristol.ac.uk (S. Moazzeni).
[1] IEEE Fellow.

https://doi.org/10.1016/j.comnet.2024.110202
Received 8 October 2023; Received in revised form 20 December 2023; Accepted 19 January 2024
Available online 30 January 2024
1389-1286/

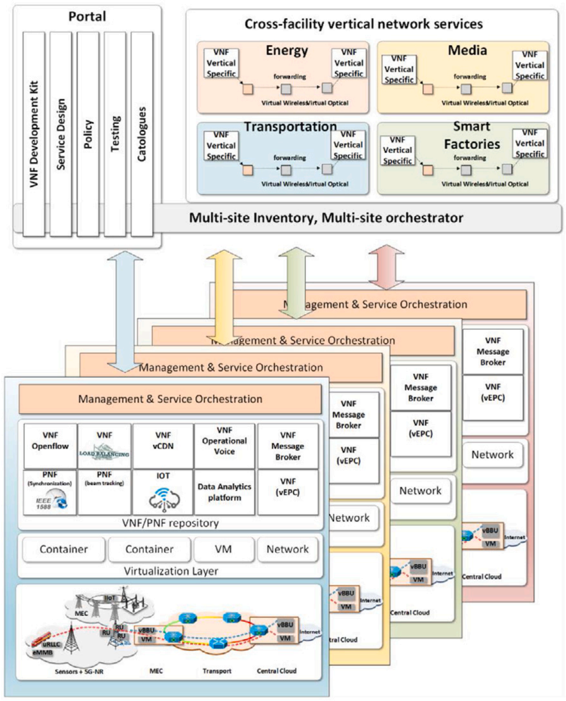

Fig. 1. 5G-VICTORI functional architecture.

inter-site connectivity, inter-domain orchestration, and support onboarding of inter-domain services as well as end-to-end slice monitoring and management for the deployed end-to-end services. This thin orchestration layer is referred to as 5G-Victori Infrastructure Operating System (5G-VIOS) responsible for the Life Cycle Managements (LCMs)of NSaaS. The selected modules of the 5G-VIOS platform enable the implementation of an intelligent inter-domain orchestrator to deploy network services across domains and provide the capability to efficiently exploit and interact with the available underlying technologies such as MANO platform, SDN controller, and the monitoring platform available at different facility sites. 5G-VIOS architecture has adopted a cloud-native and microservice-based design, allowing individual components to be developed and extended in parallel, providing flexibility and adaptability. All the modules work in unison to deliver efficient management of NSaaS to the user. Currently, docker containers for each 5G-VIOS module and a user manual with information for installation and automated deployment are publicly available [3].

The current MANO systems and recent advancements in Network Service (NS) orchestration have made significant progress in managing and profiling NSs. However, they lack the capability to autonomously and simultaneously support key functionalities such as E2E autonomous orchestration and management of NSs across multiple domains, supporting seamless connectivity, and migration of NSs in case of mobility of the users, finding a deep knowledge about the NS profiles by using profiling models, and optimising the resources assigned to NSs using ML techniques to prevent unnecessary migration of NSs. In this paper, we provide a high-level overview of the components within the 5G-VIOS architecture and to address the mentioned gaps we detail how the 5G-VIOS orchestrates and manages the NSs across multi domains. We also describe how the profiling models assist the 5G-VIOS to select the optimum amount of resources and prevent the unnecessary migration of NSs due to the lack of required resources.

The rest of the paper is organised as follows. Section 2 presents background and state of the art review on the related work. The detailed description of the proposed 5G-VIOS platform is presented in Section 3 followed by the description of network slice management in Section 4. In Section 5, experimental setup and the deployment of 5G-VIOS are explained, followed by our experimental evaluation results discussed in Section 6 while the conclusion and future works are described in Section 7.

Table 1
SoA projects targeting cloud-native cross-domain orchestration.

	Mobility support	Business automation	Cross-domain monitoring	Profiling
5G-VICTORI	✓	✓	✓	✓
5GUK Exchange [7]		✓		
5G-VINNI [8]		✓	✓	
5G-EVE		✓	✓	
INSPIRE 5GPlus [9]		✓	✓	
5GZORRO [10]		✓	✓	✓

2. Related works

For the LCM and orchestration of NSs, a reference architectural framework has been developed by the ETSI ISG NFV, commonly known as the ETSI NFV Management and Orchestration (ETSI NFV MANO) [4]. Two popular orchestration platforms, Open Source MANO (OSM) by European Telecommunications Standards Institute (ETSI) and Open Network Automation Platform (ONAP) by the Linux Foundation, are fully compliant to the ETSI NFV MANO framework and support multi-technological domains (i.e., SDN and NFV). Although the aforementioned orchestration frameworks were initially designed for the LCM of a NS and a Virtualised Network Function (VNF), their components were extended to support the LCM of a network slice. Network Slice Management Function (NSMF) and Network Slice Subnet Management Function (NSSMF) have been implemented by both solutions (OSM, ONAP) based on 3GPP TR28.801 specification [5]. These orchestration frameworks can efficiently deploy and manage the network slice within a single administrative domain. However, they still lack flexibility while managing network slices across different administrative domains [6]. This flexibility relates to the way these platforms orchestrate the underlay network infrastructure. Existing solutions mostly focus on deploying network services and do little or nothing to configuring networking, particularly for inter-domain network configurations. This limitation does not allow to create complex/adaptive inter-domain services and thus network slices.

The establishment of a multi-domain network slice instance leverages the benefits of recursive virtualisation. In this work, a domain is represented by an edge infrastructure/facility which utilises a MANO system to deploy, connect and manage the NSs. The challenge for inter-edge orchestration is not only limited to interconnecting the various edges that are geographically located apart over a secure Wide Area Network (WAN), but also interacting with the different administrative domain technologies such as local orchestrator, monitoring platform, SDN controller, etc. Several State of the Art (SoA) projects have extended these tools or have been developing similar MANO systems to fit their use case needs and also support multi-domain orchestration. However, their solutions lack support for mobility and optimal resource deployment. A comparison of 5G-VIOS's features against other platforms and projects targeting cloud-native cross-domain orchestration with a separate governance domain that manages the different administrative domains/facilities is presented in Table 1.

For example, 5GUK Exchange [7] creates an inter-domain orchestration brokering solution built upon the OSM. Additionally, it incorporates dynamic service-based L2 cross-site connectivity capabilities. However it does not support the seamless connectivity while users move or assigning optimum resources to the NSs. 5G-VINNI [8] and 5G-EVE [11] both perform inter-domain NS deployment but without the integration of a monitoring data consumer for ML-based profiling of NS performance. INSPIRE5GPlus [9] uses the ETSI Zero-touch network

Table 2
List of acronyms.

Acronym	Full form
NSSMF	Network Subnet Slice Management Function
NSMF	Network Slice Management Function
SMF	Slice Management Function
NSaaS	Network Slice as a Service
PRO	Profiling
MON	Monitoring
SBR	Service Broker
SMA	Service Manager
MOB	Mobility Manager
ICM	Inter-edge Connectivity Manager
EPA	Edge Proxy
REP	Repository
SCO	Service Composer
GUI	Graphical User Interface
AGA	API Gateway
NWDAF	Network Data Analytics Function
NEF	Network Exposure Function
NRF	Extended Network Repository Function
AF	Application Function
SEPP	Secure Edge Protection Proxy
CAPIF	Common API Framework

and Service Management (ZSM) reference architecture for trustworthy interconnection of different administrative domains, nevertheless does not focus on ML-based profiling and mobility support. Finally, 5GZORRO [10] focuses on the implementation of a zero-touch security and trust implementation again without proposing a service migration workflow.

The key innovations of 5G-VIOS are focused on mobility management and providing seamless connectivity while users move, optimising the resource allocation through profiling models, as well as extending inter-domain orchestration with L3 connections. In this work, we define 'experiments' as network slices with predefined requirements that are used to construct end-to-end services. Fig. 2 portrays the high-level functional architecture of the 5G-VIOS framework connecting two reference facilities. The design principle of 5G-VIOS follows a microservice-based architecture in which components communicate with each other over a common service bus using open Application Programming Interfaces (APIs). Through a secure edge proxy and authentication/authorisation mechanisms, i.e. Secure Edge Protection Proxy (SEPP) and Common API Framework (CAPIF), each facility owner has full control over the exposed information and the rules associated with the API access for untrusted applications. The platform acts as a trusted neutral mediator between facility owners, service developers, and Vertical users minimising contractual, legal, and communication overheads.

3. 5G-VICTORI infrastructure operating system (5G-VIOS)

The end to end reference architecture of 5G-VIOS, as shown in Fig. 2, closely resembles that of the 5G Service Based Architecture (SBA). In this architecture, the 5G-VIOS components are interconnected in a manner similar to the 5G SBA Network Functions (NFs). Detailed information about this E2E reference architecture can be found in [13]. In essence, the 5G-VIOS components extend specific 5G NFs, including the SEPP, Network Exposure Function (NEF), and Extended Network Repository Function (NRF), to facilitate inter-domain connectivity and services. This implies that functionalities of certain 5G NFs are integrated into the 5G-VIOS components. More specifically, the components defined by 3GPP are mapped to corresponding 5G-VIOS components, as detailed in corresponding functional blocks below. In addition, the list of acronyms used in the 5G-VIOS description are provided in Table 2.

- **Portal (GUI)**: experiment owners/users manage, deploy and visually monitor experiments and their performance via a web-based portal.
- **Edge Proxy (EPA)**: an extension of 3GPP's SEPP [14], Edge Proxy (EPA) facilitates secure communication between 5G-VIOS and edge facilities. It supports functions such as access token generation, bootstrapping, NF discovery, and NF management. EPA employs CAPIF to expose facility capabilities to the common **Application Function (AF)** repository in 5G-VIOS, ensuring secure communication with edge orchestrators and offering Northbound APIs for various edge components. Communication between the 5G-VIOS administrative domain and the Edge Proxy(s) is permitted through the **API Gateway (AGA)**.
- **Repository (REP)**: 5G-VIOS enhances the NRF of the SBA 5GC [15] by deploying a common repository to gather NF information and supporting inter-domain NSs. Local Extended NRFs at each facility discover available NF instances, record Virtual Network Service Descriptors (VNSDs), and maintain information about NSs such as NS profiles, EPAs, and VNSDs for edge orchestrators.
- **Service Composer (SCO)** composes inter-edge VNSDs based on user-selected NSs for deployment on edge facilities. It creates inter-domain network service descriptor (iNSD) templates and defines transport subnet slices for network slices.
- **Service Broker (SBR)** acts as an intermediary between the EPA and the rest of the 5G-VIOS components. It checks the resource availability at each edge with the help of the Profiling system and NEF. NEF provides user-friendly tools for exposing 3GPP network services and capabilities via northbound RESTful APIs. This implementation can be scaled to multiple domains through the Service Broker (SBR) component of 5G-VIOS, which exposes different edge capabilities and services in the common infrastructure (5G-VIOS) and instantiates inter-domain services through the Service Composer (SCO) and other 5G-VIOS components.
- **Service Manager (SMA)** manages the lifecycle of inter-edge network slices and includes functions of NSMF and NSSMF [5]. It deploys required NSs, establishes transport networks, monitors NSs, and manages slice instances.
- **Inter-edge Connectivity Manager (ICM)** deploys transport networks to enable secure communication between services on different edges. It serves as a bootstrapping point, connecting edges to 5G-VIOS, and establishes data paths for inter-domain network slice instances.
- **Mobility Manager (MOB)** ensures service continuity during NS migration between edges. It efficiently migrates NSs to target edges, terminates redundant services, and instructs ICM to adjust the transport network.
- **Profiling (PRO)**: a Profiling instance is integrated into each facility, allowing it to generate the performance records of NSs, profile them, and manage them by using the prediction models. As the profiling entity of the 5G-VIOS, we apply NAP [16] which is an autonomous VNF profiling method, considering multiple resource types (e.g., CPU, memory, and network) at the same time. There are some other profiling solutions for VNFs in the literature, such as [17,18], and [19] however, these existing solutions for NFV MANO do not automate the entire process (from benchmarking to data analysis) of profiling. The authors in [20,21] propose some methods for profiling the whole service function chain performance behaviour. However, their automation capability is limited, and they do not consider the profiling results of each constituent VNF individually.
  Computational profiles are created for the corresponding VNFs by utilising the collected monitoring information in combination with the selected resource configuration weights. More precisely and being a time-sensitive process, the Profiling weighted randomly selects the most impacted configuration of resources using the Weighted Resources Configuration Selection (WRCS) algorithm [16]. The edge Profiler generates the "performance dataset records" of the VNFs and NSs at each facility before the central Profiler within 5G-VIOS creates a performance model for the

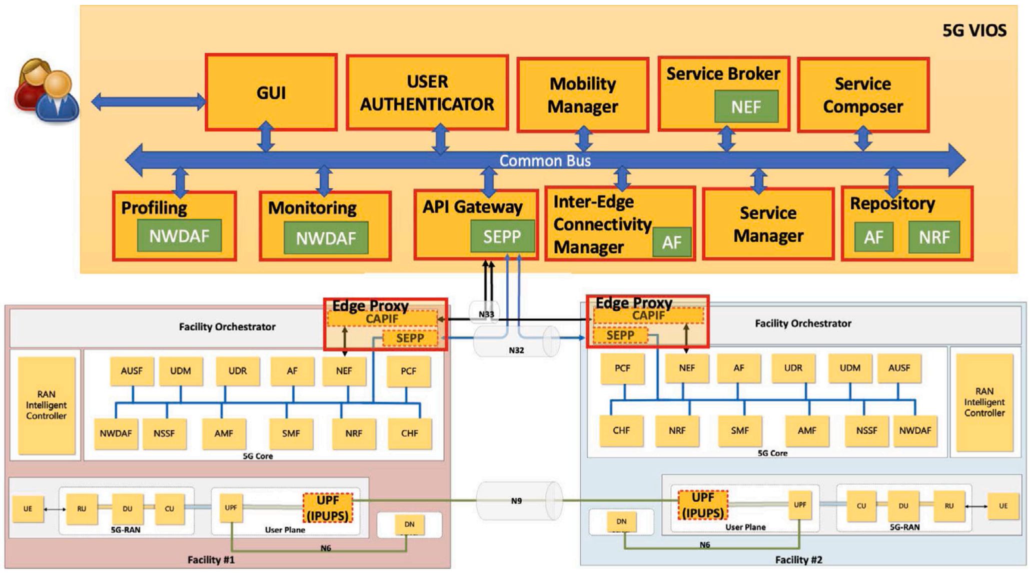

Fig. 2. An overview of the 5G-VIOS architecture [12].

inter-edge NSs and predicts the optimum network configuration and resource allocation (e.g., CPU cores, RAM, storage, and link capacity).

Through the usage of Profiling, the 5G-VIOS extends Network Data Analytics Function (NWDAF) to inter-domain architectures, enhancing performance analysis; it collects and analyses monitoring data from various sources, utilises ML techniques to predict network configuration and resource allocation for NSs, and facilitates mobility services and auto-scaling based on Quality-of-Service (QoS) measurements.

- **Monitoring System (MON)**: as shown in Fig. 3, the Monitoring (MON) system has two entities; one runs locally at each edge facility and collects metrics from different components, while the other runs within 5G-VIOS and collects the aggregated inter-domain data. The monitoring data is differentiated for each network slice, thus providing slice-specific information within a multi-slice environment. MON pulls monitoring information directly from the underlying NFV Orchestrator (NFVO), applications services, infrastructure, and monitoring tools such as the Prometheus stack [22]. Data visualisation is then made available in the form of performance figures through a dashboard in the 5G-VIOS Portal. In addition, the monitoring data is ingested internally by the Profiling system, which measures the KPIs attained for the requested services and generates performance profiles. The monitoring metrics as well as the Performance dataset records are stored using the Elastic stack (Elasticsearch, Logstash, etc.) data repository [23]. Please refer to [13] for more details on various telemetry components at the edges.

4. Network slice as a service workflow

5G-VIOS offers an easy and automated way to deploy and manage network slices across multiple domains. 5G-VIOS micro-services are linked over a common service bus delivering the functionality illustrated in Fig. 4.

Different scenarios with respect to the deployment location of the AF can be implemented depending on the capabilities of the facility controlling the Vertical application, as well as the server hosting the application (i.e. the Application Server (AS)). The role of AF in the provisioning of services with enhanced Quality-of-Experience (QoE) is critical. The feedback received from the application can influence the end-to-end network configuration and resource allocation, including the traffic routing and steering decisions of the Subnet Management Function (SMF), the selection of the Mobile Edge Computing (MEC) platform hosting the AF, the adaptation of the trigger rate, the exposure of statistics to the analytics function, etc. In multi-domain scenarios, CAPIF enables the inter-edge functionality of the AF.

5G-VIOS can facilitate the public–private network connectivity and orchestration. For example, the platform can be deployed within a trusted facility (i.e., public operator) and provide services to a non-trusted facility (private operator). To support low-latency services, each facility can host any vertical application, e.g. “AppX”, on a local MEC platform and have it controlled by the AS functionality within the 5G-VIOS. Network metrics and statistics from each facility are exposed to the monitoring and network analytics modules of 5G-VIOS through SEPP. Based on the collected metrics (including user information and location) and the application-specific performance targets, the AF affects traffic routing and steering, as well as the MEC selection process.

The NSaaS workflow is initiated by a vertical end-user, who can request a new network slice/service through the Portal (GUI). The 5G-VIOS Portal component provides the front-end for users to create, manage, and monitor their network slices. The end-users can create an experiment by selecting a set of NS packages available at the distinct edges (as advertised in the common repository), along with the required performance targets and QoS used by the Profiler. This information is compiled into an iNSD and used to deploy the NS across different facilities.

The deployment of an inter-domain network slice is done in two phases:

1. For the initiate phase, 5G-VIOS performs a resource check at each deployment edge. Then, a Profiling request and a resource prediction are made, and the orchestrators at the corresponding edges are signalled with the relevant resource allocations and descriptor checks.
2. Once successful, SMA starts the network slice deployment or instantiation phase. The composed iNSD is forwarded to the SMA component and is parsed into a template. The SMA identifies all required NSs and sends a POST request to the SBR micro-service, which is then forwarded to the EPA at each edge/facility using SOL005 APIs. EPA request is processed, and an instantiation request is sent to the edge orchestrator (e.g. OSM), which then instantiates the NS at the requested virtual infrastructure

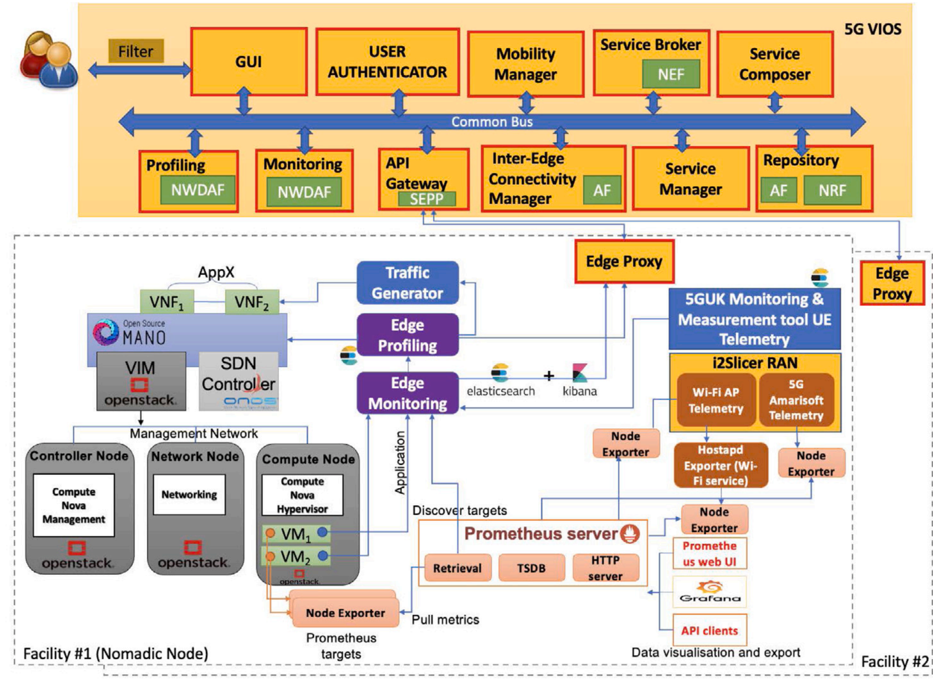

Fig. 3. E2E Monitoring at 5G-VIOS.

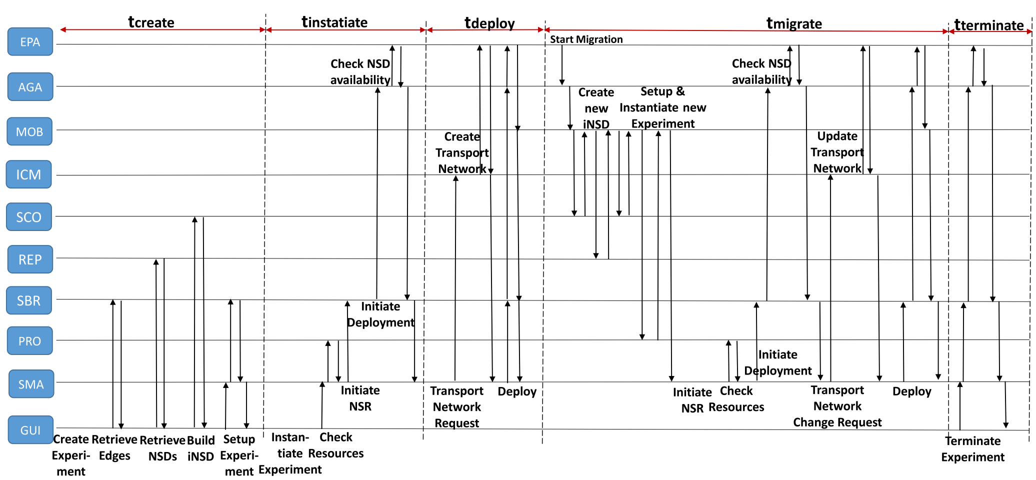

Fig. 4. 5G-VIOS workflows for creation, instantiation, deployment, migration, and termination of an experiment.

(i.e., OpenStack or Kubernetes). A VLAN is allocated to the deployed NS, allowing intra-slice communication within the edge. Once all the services of the network slice are instantiated, the ICM is instructed to create the corresponding transport network and subnet, which is based on the DMVPN architecture.

One key innovation target for 5G-VIOS is the introduction of mobility management within the LCM of the NS. This would offer seamless service and QoS continuity to vertical users while they move. When an AppX requests a service migration, the MOB makes requests to recompose the iNSD with the new target edge of the NS, after coordinating with the corresponding edge orchestrator. The new recomposed iNSD is then forwarded to the SMA, which requests a new NS deployment to the target EPA. Once the instantiation of the NS is successful, SMA terminates the previous NS instance from the current edge by sending a request to the EPA. If the end-user no longer requires the service, the network slice can be terminated through the Portal, releasing all related computational and network resources.

At the level of individual facilities, other functional entities may coexist with the edge orchestrators to assist in the intra-domain management of network slices. Such is the case of the Nomadic Node edge at the Bristol facility as one of the edges orchestrated by 5G-VIOS in 5G-VICTORI [24]. In particular, within this facility, the i2Slicer has been introduced for the management of network slices that span from the core and edge segments towards the Radio Access Network (RAN) while integrating multiple Radio Access Technologies (RAT) [25]. The primary objective of i2Slicer is to deliver both multi-tenancy and multi-service capabilities within the 5G network framework [26]. While achieving multi-tenancy can be realised through techniques like MOCN (Multiple Operator Core Network) and RAN sharing, a robust support system for multi-service functionalities necessitates a disaggregated deployment approach that leverages 5G network slicing [27]. This approach ensures dynamic and efficient resource management in alignment with the status and requirements of the various services offered. As depicted in Fig. 5, our designed disaggregated slicing approach comprises two fundamental components: a common control plane with

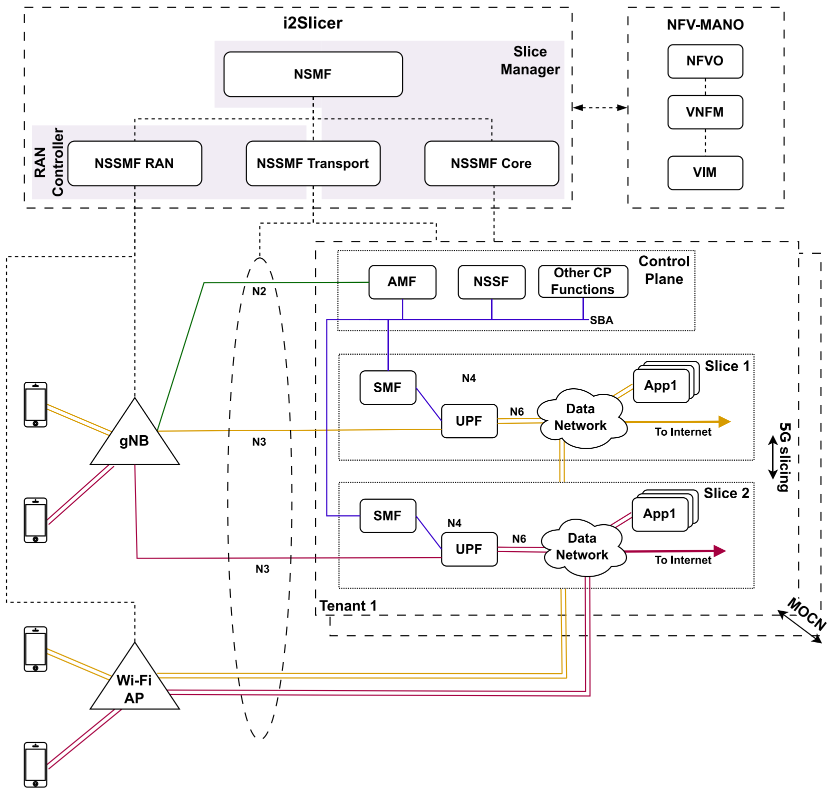

Fig. 5. i2Slicer: Multi-RAT network slicing architecture.

shared NFs serving all slices under a common operator or tenant, and an isolated user plane with dedicated SMF (Session Management Function) and UPF (User Plane Function) instances for each slice.

It is essential to highlight that our system supports multi-tenancy through MOCN, achieved by deploying additional control planes and dedicated data plane slices when required. Additionally, as illustrated in Fig. 5, it supports the deployment of multi-RAT slices by integrating the data planes of Wi-Fi and 5G RAN. The deployed multi-RAT slices are monitored through custom Prometheus Exporters, and the generated metrics are integrated into the MON system, as depicted in Fig. 3.

After the activation of slices within the Bristol facility, users can initiate requests via 5G-VIOS for the deployment of their vertical applications. These applications are seamlessly interconnected with the dedicated VLAN associated with the respective slice. This setup guarantees swift and low-latency access to these services, accessible from devices connected via both Wi-Fi and the 5G RAN.

5. Experimental setup

Fig. 6 describes the experiment setup. 5G-VIOS is used to deploy applications' inter-edge NSs at three stationary edges (Smart Internet Lab, WTC, MShed) and a non-stationary edge (Nomadic Node), providing the necessary intra- and inter-domain management and orchestration to enable end-to-end service provision and mobility. The Nomadic Node is a combination of networking (5G RAN/Core, Wi-Fi 6 AP, switches, customer-provided equipment (CPEs)), computational (servers), and software components as shown in Table 3 that provides a portable small-scale capability and services of a fully operational e2e network, including the cloud, the core, and the RAN. MEC capability is also available, which allows the instantiation of the different application services locally. The Nomadic Node backhaul connectivity to the network is delivered via multiple wireless links using a multi-modem CPE over trusted (private) or untrusted (public) 4G/5G networks. The i2Slicer at the Nomadic Node facilitates the management of network slices, extending from the core and edge segments to the RAN and seamlessly integrating various RATs.

6. Performance evaluation

We measure the average time taken 5G-VIOS to complete the process over a number of iterations considering five stages including *Experiment Creation*, *Service Instantiation*, *Service Deployment*, *Service Termination*, and *Service Migration*. Each stage of the service life cycle shown in Fig. 4 has been measured over four iterations and presented in Table 4. It can be seen that the creation (t_{create}) and instantiation ($t_{instantiate}$) phases are two orders of magnitude faster than the deployment (t_{deploy}) and termination ($t_{terminate}$) phases and three orders of magnitude faster than the migration ($t_{migrate}$) phase. This shows that the workflows related to managing the VM instances with OpenStack (deployment, migration, termination) are more time-consuming while workflows that only require 5G-VIOS or the setup of the inter-domain transport network are faster.

Regarding the intra-domain slice deployment at Nomadic Node, Table 5 illustrates the average execution time (over 100 iterations) for creating, activating, and deploying (i.e., creation plus activation) a slice using i2Slicer and considering the monolithic and disaggregated deployment modes (i.e., the first slice, which includes control and user plane NFs, and subsequent slices, which only include user plane NFs). Note that in all the evaluated cases the time spent during slice deployment was below 60 s, which is within reasonable limits for scenarios such as the temporary or pop-up network demonstrated in the Nomadic Node edge illustrated in Fig. 6. Additionally, although the disaggregated approach entails slightly longer deployment times due to the additional creation of logical resource entities and NFs, we can consider it a reasonable trade-off for its higher flexibility and efficiency in resource allocation.

Furthermore, we evaluate the performance of various ML-based prediction models within the profiling system in predicting the resource configuration and resource utilisation of a VNF. Then, we showcase and discuss the efficiency of these predictions in the proactive life-cycle management of VNFs. In order to test the performance of our profiling method, the 5G-VIOS requested the OSM at each edge to

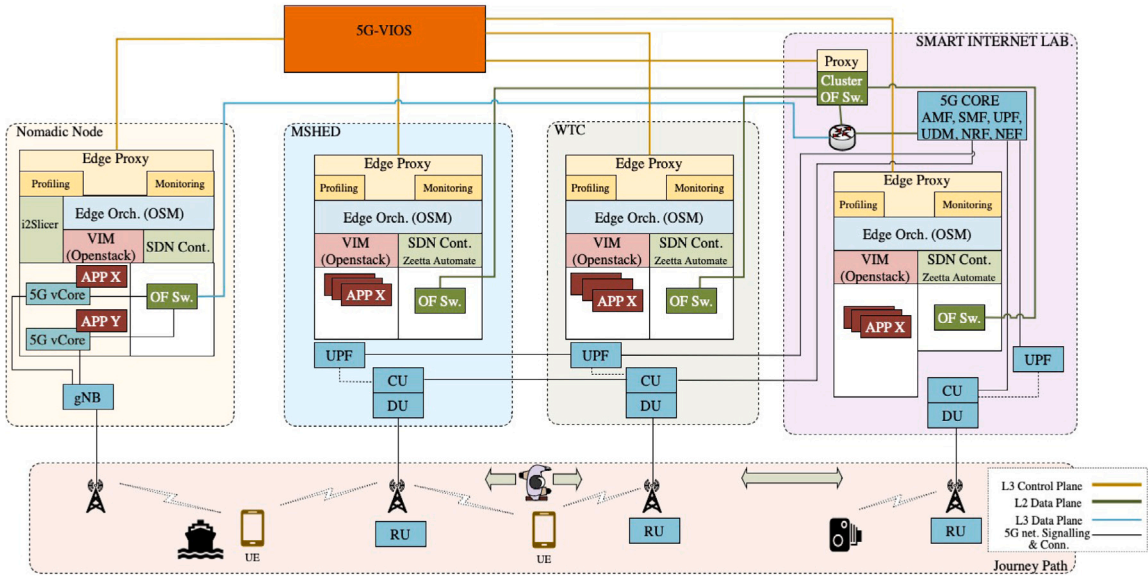

Fig. 6. Experiment setup.

Table 3
Software components of experimental setup.

Component name	Version	Notes
5GVIOS	0.1.x	Minimum resource requirements: vCPU:6 cores, RAM:16 GB, Storage:100 GB
Openstack [28]	Wallaby	Deployed within an Edge
ETSI OSM [29]	13	Deployed with EPA within an Edge as a Kubernetes cluster vCPU:5 cores, RAM:12 GB, Storage:100 GB (with EPA) vCPU:3 cores, RAM:6 GB, Storage:50 GB (without EPA)
VyOS router [30]	1.4-rolling	Supports inter-domain network connectivity minimum requirements: vCPU:1 core, RAM: 1 GB, Storage: 4 GB
Deployed NS VMs	N/A	Requirements: maximum vCPU: 2 cores, RAM: 4 GB, and Storage: 50 GB (Profiler predicts the optimum resources needed per each NS)
Amarisoft gNB	2023–06–10	Amarisoft Callbox Pro
Open5GS	v2.6.2	Monolithic: vCPU:4 cores, RAM:4 GB, Storage: 60 GB Disaggregated (First Slice): vCPU: 4 cores, RAM: 6 GB, Storage: 60 GB Disaggregated (Subsequent Slices): vCPU: 3 cores, RAM: 2 GB, Storage: 60 GB

Table 4
Performance of 5G-VIOS during service life cycle stages (as shown in Fig. 4).

	Creation (t_{create})	Instantiation ($t_{instantiate}$)	Deployment (t_{deploy})	Migration ($t_{migrate}$)	Termination ($t_{terminate}$)
Duration	0.5 s	0.62 s	66.62 s	105.81 s	18.2 s

Table 5
Time spent in seconds during slice deployment operations by i2Slicer.

Deployment modes	Slice creation	Slice activation	Slice deployment
Monolithic	13.41	25.67	**39.09**
Disaggregated: First slice	17.27	40.75	**58.02**
Disaggregated: Subsequent slices	14.36	32.71	**47.07**

deploy three VNFs (two hosting iPerfs for transmission and reception of traffic and a Firewall function in the middle as the profiled VNF instance) over OpenStack. We collected the performance dataset records of the Firewall function by exploiting the ZSM-oriented profiling entity “NAP” [16]. The collected VNF performance dataset is used to train and test the ML-based prediction models of the Profiler. To put the values of the VNF performance dataset in a scaled range, we applied a standard scaler to the dataset and then split it into a train and test dataset.

To ensure a comprehensive evaluation, we utilised the tenfold cross-validation technique on our dataset. This technique partitions the VNF performance data into ten distinct sets, each constituting one-tenth of the data. Each model is subsequently trained on nine of these sets and validated on the remaining one. This process is reiterated ten times, with the results averaged to yield a singular predictive performance metric for each model.

6.1. Profile-based resource configuration prediction

The ML approaches supported by the Profiler are *Gradient Boosting, Random Forest, Ridge Regression, K-Nearest Neighbours,* and *Neural Networks*. We exploit the regression model of these ML approaches for predicting resource configuration patterns and their classification model counterparts for predicting the resource utilisation of the profiled VNF. For the Neural Network regression and classification models, we used the Multi-Layer Perceptron (MLP) Regressor and MLP Classifier, respectively. It is worth mentioning that all the classification models for profile-based resource utilisation predictions are generic out-of-the-box SciKitLearn implementations, and we did not adjust any parameters using any hyperparameter optimisation techniques.

Table 6
Performance of ML regression models for profile-based resource configuration prediction.

ML	CPU				Memory				Link			
regression models	MAE	RMSE	MAPE	EVS	MAE	RMSE	MAPE	EVS	MAE	RMSE	MAPE	EVS
Gradient Boosting	0.03	0.05	5.57	0.93	63.61	85.97	4.87	0.79	14.81	41.60	2.60	0.89
Random Forest	0.03	0.06	5.08	0.93	58.93	82.09	4.48	0.81	14.23	42.73	2.55	0.88
Decision Tree	0.03	0.07	4.64	0.90	66.31	100.54	5.01	0.71	15.43	47.20	2.65	0.85
Ridge Regression	0.05	0.07	10.02	0.88	77.06	113.24	5.95	0.62	46.64	64.63	8.43	0.72
K-Nearest Neighbours	0.05	0.08	9.28	0.86	69.00	95.10	5.33	0.73	38.56	58.89	6.86	0.77
Neural Networks	0.04	0.06	8.13	0.92	82.68	118.89	6.33	0.59	27.08	49.58	4.96	0.84

However, to improve the accuracy and optimise the performance of the regression models, we tuned their hyperparameters using the GridSearchCV which is an exhaustive search optimisation algorithm. By exploiting the best values obtained for hyperparameters, we configured our regression models as follows: Gradient Boosting with a learning rate of 0.05, a maximum depth of 5, and 100 estimators; Random Forest with a maximum depth of 10, a minimum sample split of 5, and 150 estimators; Decision Tree with a maximum depth of 10 and a minimum sample split of 10; Ridge Regression with an alpha value of 10; K-Nearest Neighbours utilising seven neighbours and distance-based weights; and a Neural Network featuring a ReLU activation function, hidden layer sizes of (50, 50), and a learning rate of 0.01.

To evaluate the profile-based resource configuration predictions made by the mentioned regression models, we used the following evaluation metrics: (i) *Mean Absolute Error (MAE)*: measures the average magnitude of errors between predicted and actual values, providing a straightforward indication of model accuracy. (ii) *Root Mean Squared Error (RMSE)*: represents the square root of the mean of squared differences between predicted and actual values, emphasising larger errors. (iii) *Mean Absolute Percentage Error (MAPE)*: calculates the average percentage difference between predicted and actual values, enabling assessment based on relative errors. (iv) *Explained Variance Score (EVS)*: indicates the proportion of variance in the dataset captured by the model predictions, ranging from zero to one, where one denotes a perfect prediction. These different evaluation metrics help quantify the magnitude of errors in predicted resource configurations, offering insights into absolute accuracy and identifying large deviations between predicted and actual values.

The CPU resource configuration prediction results, shown in Table 6, illustrate that the Gradient Boosting and Random Forest models exhibit superior performance in predicting CPU configurations. These models achieve an EVS close to 0.93, indicating that they account for approximately 93% of the variance in the test data. The Decision Tree, although slightly lagging behind Gradient Boosting and Random Forest, still offers decent accuracy, suggesting that the relationship between the features and CPU configuration may be inherently hierarchical. While Ridge Regression and k-Nearest Neighbours have relatively higher error metrics (especially MAPE), they still offer reasonable accuracy, making them viable alternatives when computational simplicity is desired.

In the case of Memory resource configuration prediction, as shown in Table 6, Gradient Boosting and Random Forest remain consistent top performers across different resource configurations, achieving EVS scores close to 0.80 for memory predictions. The Decision Tree offers comparable accuracy, suggesting that for memory configurations, simpler models can be as effective as more complex models like Gradient Boosting, especially in scenarios where a slight trade-off in accuracy can be afforded. As with the other configurations, Gradient Boosting and Random Forest outperform other models in predicting link configurations, as illustrated in Table 6.

6.2. Profile-based resource utilisation prediction

The VNF resource utilisation percentages are categorised into five qualitative states: trivial, low, medium, high, and overloaded, indicated by numbers 0 to 4, respectively. These categories represent the set of target classes that the ML-based classification models try to predict. The boundaries for the resource range are adjustable and can be defined by the service providers. Since resource utilisation is subjective and VNF-dependent, we propose that service providers set these boundaries according to (1) their resource utilisation policies and (2) the initial resource requirements defined in the VNF descriptors.

These ML models are evaluated, based on their ability to predict the VNF resource utilisation class patterns, by using the following different error measures: (i) *Accuracy*: measures the ratio of correctly predicted instances to total instances, providing an overall assessment of model correctness. (ii) *Precision*: indicates the proportion of correctly predicted positive observations among all predicted positive observations, gauging model precision. (iii) *Recall (sensitivity)*: measures the ratio of correctly predicted positive observations to actual positive observations, indicating the model's ability to identify positives. (iv) *F1 Score*: a harmonic mean of precision and recall, offering a balanced measure of a model's accuracy and consistency. These different error measures provide a comprehensive evaluation of classification models' performance and help understand the overall correctness and precision of resource utilisation class predictions.

The performance evaluation of various classification models for resource utilisation prediction of the profiled VNF is represented in Table 7. Analysing the results for CPU utilisation classification depicted in Table 7, Gradient Boosting and Neural Networks emerge as the leading models, achieving an accuracy close to 68%. The Decision Tree and Random Forest models follow closely. For Memory utilisation classification, Logistic Regression stands out, achieving an accuracy of 84%, followed closely by the Gradient Boosting and Random Forest model. The Neural Network model's performance for Memory classification is on par with its performance for CPU utilisation classification, reinforcing its consistency across different resource utilisation types. For Link utilisation prediction the Gradient Boosting model yet again takes the lead, but what is noteworthy is the remarkable performance leap of the Neural Network model, achieving an accuracy of 93%. This suggests that the Link utilisation patterns are complex and are better captured by the intricate architectures of neural networks. The Decision Tree and Random Forest models also display commendable accuracy levels, reinforcing their reliability in diverse classification scenarios.

Computational Complexity and Prediction Performance Trade-offs: Gradient Boosting and Random Forest models, while accurate, are inherently complex. They involve ensemble methods, where multiple weak learners (usually decision trees) are combined to form a robust prediction. The training time for these models can be relatively high, especially when the number of trees is increased. However, once trained, their prediction time is swift. The Decision Tree, on its own, is computationally simpler than its ensemble counterparts. It is faster to train and offers good interpretability. In cases where real-time predictions are necessary and slightly compromised accuracy is acceptable, a Decision Tree might be preferred. Support Vector Regression, Ridge Regression, and K-Nearest Neighbours models are generally less computationally intensive than ensemble methods, especially K-Nearest Neighbours, which is instance-based. However, k-Nearest Neighbours' prediction time can be high, especially with large datasets, as it computes distances between instances. Support Vector Regression

Table 7
Performance of ML classification models for profile-based resource utilisation prediction.

ML classification models	CPU				Memory				Link			
	Accuracy	Precision	Recall	F1 Score	Accuracy	Precision	Recall	F1 Score	Accuracy	Precision	Recall	F1 Score
Gradient Boosting	0.68	0.67	0.68	0.67	0.83	0.83	0.83	0.83	0.92	0.92	0.92	0.91
Random Forest	0.66	0.64	0.66	0.65	0.83	0.82	0.83	0.83	0.90	0.89	0.90	0.88
Decision Tree	0.65	0.66	0.65	0.65	0.81	0.82	0.81	0.81	0.92	0.93	0.92	0.92
Logistic Regression	0.58	0.60	0.58	0.54	0.84	0.82	0.84	0.83	0.90	0.84	0.90	0.86
K-Nearest Neighbours	0.60	0.59	0.60	0.59	0.80	0.78	0.80	0.78	0.85	0.79	0.85	0.82
Neural Networks	0.68	0.66	0.68	0.67	0.82	0.81	0.82	0.81	0.93	0.92	0.93	0.92

and Ridge Regression, being linear models, offer good trade-offs between complexity and accuracy. Neural networks can vary widely in computational complexity based on their architecture. The model used here is a simple feed-forward network, which, while more complex than linear models, is much simpler than deep networks. Training can be time-consuming, but prediction times are generally fast.

Considering the trade-off between prediction performance and computational complexity, the Decision Tree stands out. It offers competitive accuracy, especially for CPU and Memory configurations, and is computationally simpler than ensemble methods and neural networks. For scenarios where quick predictions are crucial and a slight trade-off in accuracy is acceptable, the Decision Tree might be the best choice. In conclusion, the choice of model should be driven by the specific requirements of the application. If the highest accuracy is paramount and computational resources are abundant, ensemble methods like Gradient Boosting or Random Forest are recommended. However, if there is a need to strike a balance between accuracy and computational efficiency, the Decision Tree offers a compelling choice.

6.3. Profile-based VNF migration performance

As we explained in Section 3 the developed prediction models within the profiling system are used by 5G-VIOS to set up resources for the NSs. These predictive models are valuable not just when these services are first deployed but also throughout their life cycle. The predicted resource utilisation can be used as input for resource managers and LCM algorithms such as VNF placement, VNF migration, and topology optimisation algorithms. The value of the proposed resource utilisation predictions becomes clear when the resource manager takes proactive steps in response to the predicted decrease in resources by releasing unnecessary resources to ensure resource efficiency. In this section, we discuss a practical example to demonstrate how the information gathered from our profiling system can enhance decision-making in managing VNFs.

We focus on the proactive mode, where actions are taken beforehand to prevent issues. In our case study, we are looking at a practical scenario where moving VNFs to different physical nodes is a costly process, especially when there are not enough resources available. This migration is often prompted by two situations: either the network Link is overloaded, or the physical node's CPU and Memory resources are overwhelmed. To simplify things, we assume that each virtual machine (VM) hosts just one VNF. If a VNF starts using more resources than it should, our system triggers the migration of the entire VM to another host with sufficient resources. However, VM migration is costly and can disrupt services, so it is something we want to avoid if possible. Instead of waiting for links or nodes to become overloaded, our system uses predictions to anticipate when a VM might reach its resource limit. By doing this, we can proactively increase the VM's Link capacity and resource allocation before the situation becomes critical. This scaling prevents the VM from overloading in the first place, avoiding the need for disruptive VM migrations. This approach ensures that the services remain reliable and consistent by preventing resource-related issues from arising.

In this case, we focus on a vFW-based NS, whose profile-based prediction models were introduced in the preceding section. We select the Decision Tree aggressor as the prediction model as it provides a tradeoff between prediction performance and the computational complexity we discussed before. This model predicts if the VM's resources will get overloaded based on how much resource will be utilised. We set some thresholds to decide when a VM is overloaded. For example, in this case, study we consider a VM Link-overloaded if it uses more than 92% of its available network link, and Node-overloaded if it uses more than 95% of its CPU and/or 45% of its Memory, indicated by class label 4. Otherwise, it is a normal state, indicated by class labels 0 to 3. As the main focus is on predicting the overload states, we first check the performance of the Decision Tree classification model for predicting the overload class of CPU, Memory, and Link resources through their corresponding confusion matrices, depicted in Fig. 7. As shown in that figure, the overall accuracy for predicting Link class is impressively high, at approximately 97%. All the classes have high precision and recall values, indicating the model performs very well in predicting the Link class including both overload and normal classes. The overall accuracy for predicting Memory class is approximately 67%. The model performs best when predicting class 4 (overload state) with precision and recall of 91% and 94% respectively. In the case of CPU, the overall accuracy is approximately 42%, however, it has relatively better precision and recall scores for class 4, indicating better predictions for overload state prediction.

To better elaborate on how the profile-based trained classification model is effective in reducing VM migrations, we check if the profiled NS would be overloaded during various tests where resources and data rates change randomly. We divide these tests into seven test sets, named TS1 to TS7), as illustrated in Table 8. In this table, "N/A" is used for the migration reduction percentages when it is not possible to compute the percentage when the actual number of overloads is zero. The results we got from these tests show how useful it is to have profiles for VNFs and their predictions. The number of times we correctly predicted Link or Node overloads tells us how often we need to increase the VM's network Link or CPU and Memory resources to prevent VM migrations. This proactive scaling helps us avoid having to move the whole VM around. As you can see in this table, the profile-based trained model can predict all of the Link-overloaded states in four test sets out of seven. Even in one of the test sets (TS6) it can predict both Node-overloaded and Link-overloaded states and prevent both Link- and Node-caused VM migration by 100 percent.

7. Conclusion

In this paper, we present 5G-VIOS, an autonomous inter-domain network service orchestration framework for 5G and Beyond. 5G-VIOS is able to automate the deployment of network services across multi-domain environments, providing access to functions' repositories and services available to the various facilities. It provides ML and analytical approaches through the profiling component, which aims for optimum deployment of network services. The microservice-based design of 5G-VIOS achieves a high degree of scalability, resilience, fast deployment, quick debugging, and easy maintenance. In the future, we plan

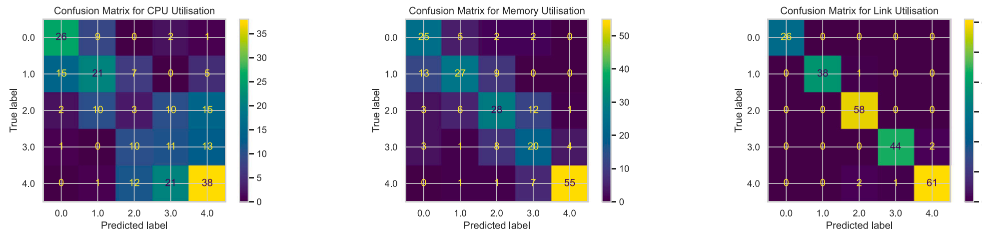

(a) (b) (c)

Fig. 7. Performance of the profile-based classification model based on the confusion matrices for predictions of (a) CPU, (b) Memory, and (c) Link resource utilisation.

Table 8
Reduction in the rate of VM migrations using profile-based Link utilisation predictions.

Test sets	Considered resource configurations			Input data rate (Mbps)	Actual *node*-overloaded states	Predicted *node*-overloaded states	Node-caused migration reduction (%)	Actual *link*-overloaded states	Predicted *link*-overloaded states	Link-caused migration reduction (%)
	vCPU cores	Memory (MB)	Link (Mbps)							
TS1	0.3	1200 to 1500	450 to 700	250 to 300	7	5	71	0	0	*N/A*
TS2	0.4	1000 to 1400	400 to 650	400 to 600	8	5	63	5	5	100
TS3	0.5	1200 to 1600	600 to 800	580 to 1060	9	4	44	6	6	100
TS4	0.6	1000 to 1300	400 to 750	750 to 1300	25	24	96	12	11	92
TS5	0.7	1000 to 1500	400 to 750	800 to 1400	14	9	64	6	6	100
TS6	0.8	1000 to 1200	400 to 700	1000 to 1400	10	10	100	6	6	100
TS7	0.9	1000 to 1400	400 to 800	900 to 1300	8	8	100	4	3	75

to develop the 5G-VIOS functionalities, such as service management, monitoring, and profiling, to support an optimised placement and orchestration of service function chains with multi-access technology requirements in 6G networks.

CRediT authorship contribution statement

Shadi Moazzeni: Methodology, Data curation, Software, Validation, Writing – original draft. **Konstantinos Katsaros:** Methodology, Writing – original draft, Software. **Nasim Ferdosian:** Methodology, Software, Formal analysis, Validation, Writing – original draft. **Konstantinos Antonakoglou:** Methodology, Software, Validation, Writing – original draft. **Mark Rouse:** Software. **Dritan Kaleshi:** Supervision, Writing – review & editing. **Adriana Fernández-Fernández:** Writing – original draft, Software, Formal analysis. **Miguel Catalan-Cid:** Writing – original draft, Software, Formal analysis. **Constantinos Vrontos:** Writing – original draft, Validation. **Reza Nejabati:** Supervision, Writing – review & editing. **Dimitra Simeonidou:** Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Shadi Moazzeni reports financial support, article publishing charges, and travel were provided by Europe Horizon 5G-VICTORI. Shadi Moazzeni reports financial support was provided by REASON under the FONRC sponsored DSIT. Nasim Ferdosian reports a relationship with Cisco-Curtin Centre for Networks that includes: employment.

Data availability

Data will be made available on request.

Acknowledgements

This work received funding from UK funded Project REASON under the FONRC sponsored DSIT, and EU project 5G-VICTORI (grant agreement No. 857201).

References

[1] S. Zhang, An overview of network slicing for 5G, IEEE Wirel. Commun. 26 (3) (2019) 111–117.
[2] 5G-VICTORI Project, Online. URL https://5g-victori-project.eu/.
[3] K. Katsaros, K. Antonakoglou, M. Rouse, 5G-VIOS user manual, 2023, http://dx.doi.org/10.5281/zenodo.8424751.
[4] NFV Release 5 Description, Standard, 2022, URL https://docbox.etsi.org/ISG/NFV/Open/.Accessed: 20-09-2023.
[5] Telecommunication Management; Study on Management and Orchestration of Network Slicing for Next Generation Network, Tech. rep., 3GPP, 2016, TR 28.801.
[6] 5G VICTORI Infrastructure Operating System (5G-VIOS) – Initial Design Specification, Deliverable d2.5, 2020.
[7] H. Alshaer, et al., The UK programmable fixed and mobile internet infrastructure: Overview, capabilities and use cases deployment, IEEE Access 8 (2020) 175398–175411, http://dx.doi.org/10.1109/ACCESS.2020.3020894.
[8] 5G-VINNI Project, 2023, Online, Last access: 3 October 2023. URL www.5g-vinni.eu/.
[9] INSPIRE5GPlus project, 2023, Online, Last access: 3 October 2023. URL https://www.inspire-5gplus.eu/.
[10] 5GZORRO project, 2023, Online, Last access: 3 October 2023. URL https://www.5gzorro.eu/.
[11] 5G-EVE Project, 2023, Online, Last access: 3 October 2023, URL https://www.5g-eve.eu/.
[12] 5G-VICTORI End-To-End Reference Architecture, Deliverable d2.4, 2022.
[13] 5G VICTORI Infrastructure Operating System (5G-VIOS) – Final Design Specification, Deliverable d2.6, 2022.
[14] G. System, Public Land Mobile Network (PLMN) Interconnection, Tech. rep., 3GPP, 2021, TS 29.573.
[15] 5G; system architecture for the 5G system (5GS), 2023, Accessed: 4-12-2023, https://www.etsi.org/deliver/etsi_ts/123500_123599/123501/16.06.00_60/ts_123501v160600p.pdf.
[16] S. Moazzeni, et al., A novel autonomous profiling method for the next-generation nfv orchestrators, IEEE Trans. Netw. Serv. Manag. 18 (1) (2021) 642–655, http://dx.doi.org/10.1109/TNSM.2020.3044707.
[17] S. Van Rossem, W. Tavernier, D. Colle, M. Pickavet, P. Demeester, Optimized sampling strategies to model the performance of virtualized network functions, J. Netw. Syst. Manage. 28 (2020) 1482–1521.
[18] S. Van Rossem, W. Tavernier, D. Colle, M. Pickavet, P. Demeester, Profile-based resource allocation for virtualized network functions, IEEE Trans. Netw. Serv. Manag. 16 (4) (2019) 1374–1388.
[19] A. Zafeiropoulos, E. Fotopoulou, M. Peuster, S. Schneider, P. Gouvas, D. Behnke, M. Müller, P.-B. Bök, P. Trakadas, P. Karkazis, et al., Benchmarking and profiling 5G verticals' applications: an industrial IoT use case, in: 2020 6th IEEE Conference on Network Softwarization (NetSoft), IEEE, 2020, pp. 310–318.

[20] S. Van Rossem, W. Tavernier, D. Colle, M. Pickavet, P. Demeester, Vnf performance modelling: From stand-alone to chained topologies, Comput. Netw. 181 (2020) 107428.
[21] M. Peuster, C. Dröge, C. Boos, H. Karl, Joint testing and profiling of microservice-based network services using TTCN-3, ICT Express 5 (2) (2019) 150–153.
[22] B. Brazil, Prometheus: Up & Running: Infrastructure and Application Performance Monitoring, " O'Reilly Media, Inc.", 2018.
[23] Elasticsearch, "Elastic stack", 2022, Accessed: 09-09-2022, https://www.elastic.co/elasticsearch/.
[24] 5G-VICTORI Final individual site facility planning, Deliverable d2.3, 2021.
[25] M. Catalan-Cid, et al., I2slicer: Enabling flexible and automated orchestration of 5G SA end-to-end network slices, in: 2023 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), 2023, pp. 159–162, http://dx.doi.org/10.1109/NFV-SDN59219.2023.10329615.
[26] A. Fernández-Fernández, et al., Validating a 5G-enabled neutral host framework in city-wide deployments, Sensors 21 (23) (2021) http://dx.doi.org/10.3390/s21238103.
[27] A. Papageorgiou, A. Fernández-Fernández, S. Siddiqui, G. Carrozzo, On 5G network slice modelling: Service-, resource-, or deployment-driven? Comput. Commun. 149 (2020) 232–240, http://dx.doi.org/10.1016/j.comcom.2019.10.024.
[28] OpenStack, Website. URL https://www.openstack.org/.
[29] Open Source MANO, Website, www.osm.etsi.org/.
[30] VyOS, Open Source Router and Firewall Platform. URL https://vyos.io/.

Shadi Moazzeni is a Research Fellow with the University of Bristol, Bristol, UK, where she is a member of the Smart Internet Lab and the High-Performance Networks research group. She is also the cluster lead researcher of the EU Horizon 2020 5G-VICTORI project and the deputy leader of WP5 in the UK funded Project REASON. She received her M.Sc. degree in Computer Architecture Engineering from Amirkabir University of Technology (Tehran Polytechnic), Tehran, Iran in 2010, and her Ph.D. in Computer Architecture Engineering at the University of Isfahan, Iran in 2018. From July 2016 to February 2017, she was a Ph.D. visiting researcher at the University of Bologna, Italy. Her current research focuses on next generation intelligent network orchestration, multi-access edge computing, 6G networks intelligent profiling and orchestration towards zero-touch network and service management.

Dr Kostas Katsaros - Head of Future Networks and 5G.

Konstantinos (Kostas) Katsaros is the Head of Technology in Future Networks and 5G at Digital Catapult, where he provides oversight and guidance on specific 5G system architecture and subsystem implementations to Digital Catapult projects, especially around new architectures using virtualization and edge computing. Prior to that, he was a Postdoctoral Research Fellow at the Institute for Communications Systems (ICS) at the University of Surrey, working on Mobile Edge Computing for Connected and Autonomous Vehicles. He holds an M.Sc. and Ph.D. from the University of Surrey on Electronic Engineering and a university degree (MEng equivalent) on Electrical and Computer Engineering from the University of Patras, Greece. His research has been focused on Connected and Autonomous Vehicles and the development of ultra-low-latency, highly reliable and secured 5th Generation (5G) Communications and Mobile Edge Computing (MEC) technologies for the Connected and Autonomous Vehicles (CAV) and Cooperative Intelligent Transport Systems (C-ITS). He has participated in several European and UK funded projects, has published several peer-reviewed papers and holds two patents.

Nasim Ferdosian is a senior research associate at Cisco-Curtin Center for Networks at Curtin University. She was a research fellow at Cergy Paris Université in 2020-2021, and a postdoctoral researcher at Dublin City University and EURECOM in 2017–2019, participating in various EU and national funded research projects. She received her Ph.D. in Computer Science from University Putra Malaysia in 2017. Her current research interests are mainly focused on network automation, data-driven network control and management, and AI and machine learning for next-generation wireless networks.

Konstantinos Antonakoglou is a Future Networks Technologist at Digital Catapult since 2021. In the past he has worked as a Research Associate at King's College London at the Centre for Telecommunications Research after completing a Ph.D. there. He has worked on EU and EPSRC funded projects such as 5GVICTORI, INITIATE, 5G-CAR and Primo-5G and is interested in a variety of topics including haptics and bilateral teleoperation over networks, inter-domain management and orchestration of network services as well as quantum communication and clock synchronisation.

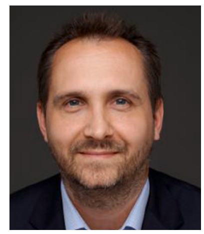

Mark is a Senior Python developer building modern web applications, RESTFul APIs, and scripts for cloud, network and process automation. He is experienced in using technologies such as Django, FastAPI, Docker, Helm, and Kubernetes in building and deploying self-contained Python applications to both AWS and OpenStack environments. He applies an Agile ethos, and DevOps best practices to automate and streamline a CI/CD workflow for rapidly building, deploying, and maintaining web applications.

Dr. Dritan Kaleshi - Director of 5G Technology.

With 30 years experience as a technologist and researcher in communication networks, distributed system design and interoperability, Dritan is Director of 5G Technology at Digital Catapult and Co-Director of SONIC Labs. Dritan has established and leads the Future Networks 5G and Digital Infrastructure programs in Digital Catapult, delivering technology development and innovation in advanced ICT systems and their adoption in enterprise networks in the UK. He was a founder of the SONIC Labs and of UK Telecommunication Network (UKTIN). He has published over 75 papers, holds three patents, and has edited two international standards on interoperability, and is often invited to speak on advanced digital infrastructure. He has represented the U.K. in international standardisation bodies and served on/chaired international conferences and technical committees.

Adriana Fernández-Fernández is a Senior Researcher at i2CAT Foundation, Spain. Currently, her main research focus is on cloud continuum and cloud-native technologies, as well as on management, orchestration, and network slicing in 5G networks. She graduated in Telecommunications and Electronic Engineering; and since 2018, she holds a Ph.D. in Telematics Engineering from the Universitat Politècnica de Catalunya (UPC). Before joining i2CAT, she worked as a postdoctoral researcher within the BAMPLA Research Group at UPC. She is an author/co-author of over 20 papers in peer-reviewed international journals and conference proceedings. She has been awarded several scholarships from the Spanish Government through the Research Training Program (FPI). Her research interests also include communication network modelling and optimisation, energy-aware routing, NFV, SDN and traffic engineering.

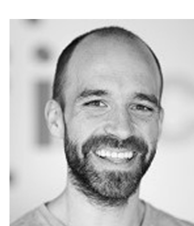

Miguel Catalan received the master's and Ph.D. degrees from the Polytechnic University of Catalonia (UPC), Spain, since 2008 and 2016, respectively. He is currently a Senior Researcher at i2CAT in Barcelona, Spain. He has an extensive experience in wireless networks, analysis and definition of communication protocols, utilisation of simulation tools, and programming embedded systems and micro-controllers. He is currently being involved in different SNS projects related to beyond 5G and O-RAN technologies.

Constantinos is a senior research associate at Smart Internet Lab, part of the 5GUK Testbed team. Constantinos obtained his Ph.D. degree in Communications in 2022 with subject "Investigating mmWave Spectrum Sharing between Mobile Operators". He finished his integrated undergraduate degree (MEng) in Electrical and Electronic Engineering at the same university in 2016. His professional role is mainly focused around the 5GUK Test Network, from the cloud to the RAN and access networks, as well as the end-user devices.

Reza Nejabati is currently a chair professor of intelligent networks and head of the HighPerformance Network Group in the Department of Electrical and Electronic Engineering in the University of Bristol, UK. He is also a visiting professor and Cisco chair in the Cisco centre for Intent Based Networking in the Curtin University, Australia. He has established successful and internationally recognised experimental research activities in "Autonomous and Quantum Networks". Building on his research, He co-founded a successful start-up company (Zeetta Networks Ltd) with 25 employees and £6 m funding. His research received the prestigious IEEE Charles Kao Award in 2016 and has done important contributions in 5G, smart city, quantum communication, and future Internet experimentation.

Dimitra Simeonidou is a full Professor with the University of Bristol, the Co-Director of the Bristol Digital Futures Institute and the Director of Smart Internet Lab. Her research is focusing in the fields of high performance networks, programmable networks, wirelessoptical convergence, 5G/6G and smart city infrastructures. She is increasingly working with social sciences on topics of digital transformation for society and businesses. She has been the Technical Architect and the CTO of the Smart City Project Bristol Is Open. She is currently leading the Bristol City/Region 5G urban pilots. She has authored and coauthored over 600 publications, numerous patents, and several major contributions to standards. She has been the co-founder of two spin-out companies, the latest being the University of Bristol VC funded spin-out Zeetta Networks, delivering SDN solutions for enterprise and emergency networks. Prof. Simeonidou is a Fellow of Royal Academy of Engineering and a Royal Society Wolfson Scholar.