This is a postprint/accepted version of the following published document:

Németh, B., et al. Delay and reliability-constrained VNF placement on mobile and volatile 5G infrastructure. In: *IEEE TRANSACTIONS ON MOBILE COMPUTING (Early access)*, Jan. 2021

DOI: <https://doi.org/10.1109/TMC.2021.3055426>

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Delay and reliability-constrained VNF placement on mobile and volatile 5G infrastructure

Balázs Németh, Nuria Molner, Jorge Martín-Pérez,
Carlos J. Bernardos, Antonio de la Oliva, and Balázs Sonkoly

Abstract—Ongoing research and industrial exploitation of SDN and NFV technologies promise higher flexibility on network automation and infrastructure optimization. Choosing the location of Virtual Network Functions is a central problem in the automation and optimization of the software-defined, virtualization-based next generation of networks such as 5G and beyond. Network services provided for autonomous vehicles, factory automation, e-health and cloud robotics often require strict delay bounds and reliability constraints influenced by the location of its composing Virtual Network Functions. Robots, vehicles and other end-devices provide significant capabilities such as actuators, sensors and local computation which are essential for some services. Moreover, these devices are continuously on the move and might lose network connection or run out of battery, which further challenge service delivery in this dynamic environment. This work tackles the mobility, and battery restrictions; as well as the temporal aspects and conflicting traits of reliable, low latency service deployment over a volatile network, where mobile compute nodes act as an extension of the cloud and edge computing infrastructure. The problem is formulated as a cost-minimizing Virtual Network Function placement optimization and an efficient heuristic is proposed. The algorithms are extensively evaluated from various aspects by simulation on detailed real-world scenarios.

Index Terms—5G, URLLC, robots, cloud, edge, VNF placement, optimization

1 INTRODUCTION

5G (& beyond) systems have been promising several appealing sometimes unbelievable future use cases and applications which could reshape our society. The vast number of IoT devices, autonomous vehicles and different types of robots collaborating with each other and with humans are expected to be part of our lives. These devices usually require coordination or fine granular, dynamically programmable control from a reliable and permanently available platform. Coordination of collaborating robots, drone swarms, self-driving cars or any types of unmanned vehicles are good examples with several fields of application from industry to agriculture and from logistics to emergency

management. The envisioned use cases typically pose serious challenges on the underlying networks and cloud platforms in terms of latency and reliability. For example, 3GPP specified a dedicated set of features for mission critical applications referred to as Ultra-Reliable Low-Latency Communication (URLLC) [1].

Edge and fog computing, multi-access edge computing (MEC) are key enablers of these applications. The main concept is to extend traditional cloud computing by deploying compute resources closer to customers and end devices. By these means, both end devices and central cloud servers can offload computational tasks to resources at the edge or the fog resulting in lower delays and in reduced network load. In order to meet the strict delay and reliability requirements of mission critical applications, a distributed and heterogeneous infrastructure and the encompassed compute and network resources should be managed carefully. The underlying infrastructure includes both public and private cloud/edge resources [2] providing execution environments for Virtual Network Functions (VNFs) interconnected by public 5G networks and privately operated domains. In this environment, resource orchestration is a challenging task which aims at always finding the proper placement of software components realizing the service. Moreover, robots or different vehicles equipped with sensors, actuators and local computation environments, provide capabilities which can or must be consumed by certain applications. More exactly, now we can run VNFs on these continuously moving mobile devices, and the uninterrupted communication to other service components should also be guaranteed. Beside mobility, the limited battery capacity and the VNFs power consumption are novel aspects to be considered in the placement decision.

This paper addresses the temporal aspects and conflicting traits of reliable, low latency service deployment over a network where mobile compute nodes (e.g., robots, drones) act as an extension of the cloud and edge computing infrastructure.

The research contribution is threefold. First, the VNF placement problem is formulated as a cost-minimizing optimization problem. The present work extends formulations in the state of the art imposing the radio coverage of mobile fog devices, and preventing that VNF deployments use fog devices that may run out of battery. Second, the optimization problem is solved by a novel heuristic algo-

- Balázs Németh and Balázs Sonkoly are with MTA-BME Network Softwarization Research Group, Hungary Emails: balazs.nemeth@tmit.bme.hu, sonkoly@tmit.bme.hu
- Nuria Molner is with IMDEA Networks Institute, Spain Email: nuria.molner@imdea.org
- Nuria Molner, Jorge Martín-Pérez, Carlos J. Bernardos, and Antonio de la Oliva are with Universidad Carlos III de Madrid, Spain Emails: —, jmartinp@it.uc3m.es, cjbc@it.uc3m.es, aoliva@it.uc3m.es

Fig. 1: Deployment of a cloud robotics warehousing NS.

rithm that, to the best of the authors' knowledge, is the first one getting close to optimal results while tackling both radio coverage, and battery restrictions of fog environments. Finally, the proposed algorithms are evaluated via extensive simulations on a real-world scenario. The results confirm the beneficial properties of our solutions in terms of scalability, cost, and runtime.

The rest of the paper is organized as follows. Sec. 2 introduces a future use case motivating our work. Sec. 3 is devoted to the detailed description of our model and the optimization problem. Sec. 4 describes the proposed heuristic algorithm. Sec. 5 presents the algorithms' evaluation from different aspects based on extensive simulations. In Sec. 6, a summary on the related work is given while Sec. 7 draws the conclusions.

2 USE CASE: MOBILE ROBOTICS

This work tackles the mobile robotics use case [3, Table 5.3.1.1-1], as a warehousing solution for future factories [4, Section 3.1.2]. In particular, it deals with the transport of goods from boats to specific locations of Valencia city haven.

The use case considers a cluster of robots, that move in a *master-slave* fashion to deliver goods arriving to the haven. Each of the robots carries containers from a pick up point (S in Fig. 1) to a drop off point (D1 and D2). In particular, the *master* robot is followed by the other *slave* robots (represented in Fig. 1) of the cluster along its way towards the drop off point. Robots communicate among themselves to report position status, or other context information useful for the *master-slave* coordination. Thus, robots have device-to-device communication between them, and computational capabilities so they can execute lightweight VNFs [5] as the *driving* and *follow* VNFs represented in Fig. 1. The *driving* VNF runs in the *master* robot to drive it towards the drop off point, and the *follow* VNFs run in the *slave* robots to follow the master robot movements until it reaches a drop off point. The *driving* VNF receives driving instructions from the *remote ctrl* VNF running in the edge server in Fig. 1, and reports sensor data like the speed to the *Database (DB)* VNF running in the cloud.

To enhance the robots' remote driving, the communication between the *remote ctrl* VNF and the *driving* VNF is crucial, indeed, resources proximity is needed as Mobile robotics demand communications with cycle times between 1ms and 100ms (for machine control, and video operated remote control cases) [3]. Thus, the placement of both *driving* and *remote ctrl* VNFs should satisfy latencies below 100ms.

While moving, robots may run out of battery or switch between RUs coverage area (see Fig. 1). Whenever the *master* robot enters a new coverage area, it attaches to a new RU to keep the connectivity with the servers running the *remote driving* VNF, and *DB* VNF (edge, and cloud servers in Fig. 1). Therefore, it is important to take into account that a robot is not selected for goods delivery if (i) it may run out of battery; or (ii) it may lose RU connectivity as it moves towards the drop off point.

To increase the RUs coverage and improve the end-to-end (e2e) delay, the use case presented in this section considers that the haven is covered by LTE RUs managed by a network operator, and New Radio (NR) RUs belonging to its Non Public Network (NPN). This is called an NPN deployment in a public network [2].

That is, Valencia city haven only owns the NR RUs, and its management (subscription, gateways, control plane) is done by the public network, i.e., a network operator.

For the public network infrastructure, a 5G transport network is assumed based on [6] and [7]. All the RUs present in the use case transmit their traffic up to an access ring composed of several switches connected in a ring fashion. The traffic of the access rings is latter gathered by the aggregation rings which forward traffic up to the core of the public infrastructure. The presented use case, assumes that cloud servers are in the core of the public network, edge servers are co-located next to the access ring and the aggregation ring switches. Regarding computational resources (i.e., CPU, memory and disk), edge servers in access rings are less powerful than edge servers in aggregation rings, and cloud servers are more powerful than edge servers.

It is worth highlighting that the problem formulation presented in this work will hold for public and private deployments, being the only consideration the cost of connec-

tion, that may vary depending on the type of management. We mention them in this work for a better understanding on the real situation in the mentioned city haven.

3 PROBLEM FORMULATION

This section presents the formulation of the use case to tackle the VNF allocation as an optimization problem. The problem is solved using an integer program solver to gain optimality and scalability insights.

3.1 System model

The network infrastructure is represented by a graph G_I , where the nodes $V(G_I)$ contain NR and LTE RUs, generally referred to as Access Points (APs) $V_{AP}(G_I)$, server nodes (representing edge or cloud servers) $V_S(G_I)$, and mobile nodes $V_M(G_I)$. Hence, the vertex set of the graph is built up as $V(G_I) = V_{AP}(G_I) \cup V_S(G_I) \cup V_M(G_I)$. Host nodes N_i with computation capacities C_{N_i} are stored in $V_H(G_I) = V_S(G_I) \cup V_M(G_I)$, and their corresponding unitary price is represented by p_{N_i} . As a realistic generalization to the mobile robotics use case, the concurrent management of multiple robot clusters is assumed. The subsets of $V_M(G_I)$ define the clusters of robots $V_{RC_q}(G_I) \subseteq V_M(G_I)$, $1 \leq q \leq Q$, where Q refers to the number of clusters. Moreover, graph edges $E(G_I)$ represent the connections between the infrastructure nodes, which are annotated by their transmission delays. Due to the mobile clusters' mobility, their connections to the static part of the infrastructure are not represented by edges in G_I .

The mobile nodes $V_M(G_I)$ are connected to access points $V_{AP}(G_I)$ in order to communicate with other nodes of the infrastructure. However, the nodes are moving and may encounter areas with overlapping access point coverage or areas where handover between different access points is needed to guarantee the connection to the servers deeper in the infrastructure. Thus, this work assumes that each AP has an associated coverage area AP_k , and the mobility pattern of robot cluster q is modeled by the probability distribution of being in the AP coverage areas $\mathbb{P}_{AP_k^q}(t)$, referred as *coverage probability* throughout the paper. Notice that a cluster can be in an area where several access points have coverage, with a different probability for each of them. Each value models the probability of a robot cluster q to fall inside the coverage area of each AP in each moment t . This model is able to compute the placement of NSs with guarantees of communication between the mobile and fixed parts of the infrastructure, while considering any model of coverage areas, such as [8] or a linear model, by using precomputed values of the coverage. The parameter t is a time instant within an interval (t_0, t_1) in which the network service will be running. For the sake of simplicity in the model, the time interval is discretized in subintervals, thus continuous time $t \in (t_0, t_1)$ becomes discrete time $t_u \in \{t_0, t_a, t_b, \dots, t_1\}$ with $t_0 \leq t_a \leq t_b \leq \dots \leq t_1$. Subintervals help to identify the moments when handovers may occur during the service time. The time division guarantees the communication between robots and APs selected in each subinterval. Note that VNFs are deployed on the same servers during all the service time, thus, they must have communication with the APs selected in each subinterval.

The cost of using an AP for a single subinterval t_u by any single cluster is p_{AP_k} . The energy consumption of the mobile nodes is modeled by the distribution $\mathbb{P}_{bat}(N_i, C_{N_i})$ depending on the allocated load to node N_i , which represents the probability of having a not depleted battery for the whole interval (t_0, t_1) . Both $\mathbb{P}_{AP_k^q}(t)$ and $\mathbb{P}_{bat}(N_i, C_{N_i})$, are used in the optimization problem to ensure robots' radio coverage, and battery needs are met during the interval (t_0, t_1) .

The requested Network Services are represented with a NS graph G_S , with the nodes being VNFs $v \in V(G_S)$ and their capacity requirements C_v . Each Service Function Chain (SFC) is a subgraph $G_s \subseteq G_S$ with its own set of VNFs and path, as the one depicted in the NS graph of Fig. 1, and expressed in Eq. (1).

$$\mathcal{C}_{SFC} = \{(G_s, \Delta_{G_s}) \mid V(G_s) \subseteq V(G_S), \\ E(G_s) \in \mathcal{P}(G_S), \Delta_{G_s} \in \mathbb{R}^+\} \quad (1)$$

where $\mathcal{C}_{SFC}$ represents the set of SFCs in Network Service G_S , and $\mathcal{P}(G_S)$ represents the paths of the NS graph G_S . Each SFC has a corresponding delay requirement Δ_{G_s} which defines an upper bound of the total delay of the SFC path $E(G_s)$.

For a better understanding of the model, all the notations used for the mathematical formulation of the optimization problem are gathered in Table 1.

3.2 Optimization problem

Our optimization problem is summarized in Formulation 1 and described in details below. The optimization must decide which infrastructure node $N_i \in V(G_I)$ should host which VNF $v \in V(G_S)$, this is represented by the binary decision variable $x(v, N_i)$ and constraints Eq. (2) and Eq. (3). The resource capacities C_{N_i} must be respected by the load allocation on each node N_i . This requirement is gathered in Eq. (4), where C_{N_i} stands for the allocated resources in infrastructure node N_i as presented in Eq. (5).

Furthermore, there may be a necessity of applying placement policies and VNF functional types. In order to include those policies in the model, the matrix $L(v, N_i)$ expresses locality constraints between the VNFs $v \in V(G_S)$ and infrastructure node $N_i \in V(G_I)$. Each element of the matrix is a binary constant, identifying whether the VNF can be located in an infrastructure node, as expressed in Eq. (6). In the use case presented in Section 2, $L(v, N_i)$ enforces the deployment of the *driving* and *follow* VNFs in the robots (i.e., mobile nodes). This requirement may be useful for other use cases, such as UAVs running virtual access points that forward traffic to the cloud (see [9], [5]). Under such scenarios, $L(v, N_i)$ can be used to enforce virtual access points to run on top of UAVs.

3.2.1 Radio coverage constraints

The deployment must also decide at each time interval to which access point each cluster of robots is attached to, that is, $AP_k^q(t_u) = 1$ in case robot cluster RC_q is connected to access point AP_k at time t_u . Eq. (7) reflects the assumption that each cluster can only be attached to one AP at each interval. The deployment decision must also ensure that the coverage probability is above the imposed threshold κ_q for mobile cluster q , representing the requirements each cluster

TABLE 1: Definition of variables and parameters

Notation		Definition
Parameters		
G_I		Network infrastructure graph
G_I	$V(G_I)$	All infrastructure nodes
	$V_*(G_I)$	Nodes of type * in the infrastructure $* \in \{AP, S, M, H, RC_q\}$
	$E(G_I)$	Edges of the infrastructure
G_S		Network service graph
G_S	$V(G_S)$	All VNFs of the network service
	$\mathcal{P}(G_S)$	All paths of the service graph
	G_s	Graph of SFC G_s
	$V(G_s)$	VNFs of SFC G_s
	$E(G_s)$	Edges of SFC path G_s
	Δ_{G_s}	Delay requirement of SFC G_s
	th_{bat}^s	Battery threshold for SFC G_s
	$\mathcal{C}_{SFC}$	Set of all SFCs
VNF v		Virtual Network Function v
VNF v	C_v	Capacity demand of VNF v
	L	Locality matrix $V(G_S) \times V(G_I)$
N_i	C_{N_i}	Total resource capacity of node N_i
	p_{N_i}	Cost per resource unit used of node N_i
	$D_{AP,S}(N_i, N_j)$	Delay between N_i and $N_j \in V_{AP}(G_I) \cup V_S(G_I)$
	$D_{M_q}(N_i, N_j)$	Delay between N_i and $N_j \in V_{RC_q}(G_I)$
	$\mathbb{P}_{bat}(N_i, C_{N_i})$	Probability of having battery for the whole time interval using C_{N_i} resources
AP_k	d_{AP_k}	Delay for the coverage area of AP_k
	$\mathbb{P}_{AP_k^q}(t_u)$	Probability of cluster q to be in the coverage area of AP_k in time subinterval t_u
	p_{AP_k}	Cost of usage of AP_k
	κ_q	Coverage probability threshold for cluster q
Variables		
$d_{G_s}(t_u)$		Delay of SFC G_s in time t_u
$d(N_i, N_j, t_u)$		Delay between nodes N_i and N_j in time t_u
$x(v, N_i)$		Placement of VNF v in node N_i
C_{N_i}		Resource usage in node N_i
$AP_k^q(t_u)$		Usage of AP_k by cluster q time t_u
$\mu : V(G_S) \mapsto V_H(G_I)$		VNF to host node mapping structure
$\alpha : \{t_u\} \times \{q\} \mapsto V_{AP}(G_I)$		AP selection structure for all clusters

$$d(N_i, N_j, t_u) = \begin{cases} \sum_{AP_k \in V_{AP}(G_I)} AP_k^q(t_u) [D_{M_q}(N_i, N^{(r_q)}) + d_{AP_k} + D_{AP,S}(AP_k, N_j)], & \text{if } N_i \in V_{RC_q}(G_I) \wedge N_j \in V_S(G_I); \\ d(N_j, N_i, t_u), & \text{if } N_j \in V_{RC_q}(G_I) \wedge N_i \in V_S(G_I); \\ D_{AP,S}(N_i, N_j), & \text{if } N_i, N_j \in V_S(G_I) \cup V_{AP}(G_I); \\ D_{M_q}(N_i, N_j), & \text{if } N_i, N_j \in V_{RC_q}(G_I); \\ \sum_{\substack{AP_{k_1}, AP_{k_2} \\ \in V_{AP}(G_I)}} AP_{k_1}^{q_i}(t_u) AP_{k_2}^{q_j}(t_u) \left(D_{AP,S}(AP_{k_1}, AP_{k_2}) + d_{AP_{k_1}} + d_{AP_{k_2}} + \sum_{n \in \{i,j\}} D_{M_{q_n}}(N_n, N^{(r_{q_n})}) \right), & \text{if } N_i \in V_{RC_{q_i}}(G_I) \wedge N_j \in V_{RC_{q_j}}(G_I) \end{cases} \quad (14)$$

needs to guarantee connectivity during the time interval, as stated in Eq. (8). Notice that Optimization problem 1 only needs to know whether cluster q has radio coverage of AP_k at time t_u . Hence, Optimization problem 1 is agnostic about how $\mathbb{P}_{AP_k^q}(t_u)$ is obtained, and the values could be derived from any radio access model. For instance, Sec. 5 obtains $\mathbb{P}_{AP_k^q}(t_u)$ with a linear function directly proportional to the distance between q and AP_k .

3.2.2 Delay constraints

In order to measure the distances between infrastructure nodes, the metric used is the delay, which in the case of the static nodes is given in a matrix containing the precomputed and the time-independent delays, $D_{AP,S}(N_i, N_j) \forall N_i, N_j \in V_S(G_I) \cup V_{AP}(G_I)$.

Similarly, the distances inside each mobile cluster are time invariant, precalculated and stored in matrix $D_{M_q}(M_i, M_j) \forall M_i, M_j \in V_{RC_q}(G_I), 1 \leq q \leq Q$.

Each access point $AP_k \in V_{AP}(G_I)$ provides

Formulation 1 Optimization problem

$$x(v, N_i) \in \{0, 1\} \quad \forall v \in V(G_S), \forall N_i \in V(G_I) \quad (2)$$

$$\sum_{N_i \in V(G_I)} x(v, N_i) = 1 \quad \forall v \in V(G_S) \quad (3)$$

$$C_{N_i} \leq \bar{C}_{N_i}, \quad \forall N_i \in V(G_I) \quad (4)$$

$$C_{N_i} = \sum_{v \in V_S(G_I)} x(v, N_i) C_v, \quad \forall N_i \in V(G_I) \quad (5)$$

$$x(v, N_i) \leq L(v, N_i), \quad \forall v \in V(G_S), \forall N_i \in V(G_I) \quad (6)$$

$$\sum_{AP_k \in V_{AP}(G_I)} AP_k^q(t_u) = 1, \quad \forall 1 \leq q \leq Q, \forall t_u \in (t_0, t_1) \quad (7)$$

$$\sum_{AP_k \in V_{AP}(G_I)} AP_k^q(t_u) \cdot \mathbb{P}_{AP_k^q}(t_u) \geq \kappa_q, \quad \forall 1 \leq q \leq Q, t_u \in (t_0, t_1) \quad (8)$$

delay equation is presented in Eq. (14)

$$d_{G_s}(t_u) = \sum_{\substack{(v_i, v_j) \in E(G_s) \\ N_i, N_j \in V(G_I)}} x(v_i, N_i) x(v_j, N_j) d(N_i, N_j, t_u) \quad (9)$$

$$d_{G_s}(t_u) \leq \Delta_{G_s}, \quad \forall (G_s, \Delta_{G_s}) \in \mathcal{C}_{SFC}, \forall t_u \in (t_0, t_1) \quad (10)$$

$$\begin{aligned} \mathbb{P}_{bat}(N_i, C_{N_i}) &= \mathbb{P}_{bat}(N_i, 0) - \\ &\quad - \frac{C_{N_i}}{\bar{C}_{N_i}} (\mathbb{P}_{bat}(N_i, 0) - \mathbb{P}_{bat}(N_i, \bar{C}_{N_i})), \quad \forall N_i \in V_M(G_I) \end{aligned} \quad (11)$$

$$\mathbb{P}_{bat}(N_i, C_{N_i}) \geq th_{bat}^s, \quad \forall N_i \in V_M(G_I), \forall G_s \in \mathcal{C}_{SFC} \quad (12)$$

$$\min \sum_{N_i \in V(G_I)} C_{N_i} \cdot p_{N_i} + \sum_{t_u, q, k} AP_k^q(t_u) \cdot p_{AP_k} \quad (13)$$

a time- and distance-independent delay to its whole coverage area, its value is denoted by d_{AP_k} , while delay between APs is given with the value, $D_{AP,S}(AP_{k_1}, AP_{k_2})$. The delay value between a mobile cluster and the static part of the infrastructure and between mobile nodes belonging to two different clusters might vary according to the assigned APs during the time interval (t_0, t_1) . A mobile cluster q has an appointed relay node $N^{(r_q)}$ (in our case the *master* robot), which is connected to the APs, and all the traffic of other mobile nodes of the same cluster towards the fixed part of the infrastructure goes through the corresponding relay mobile node. Thus, the orchestration system can execute the handover of the cluster by only connecting the relay node to a different AP. This way the delay of device-to-device communication is accounted in a different variable than the AP delays. Hence, the general delay function which covers any pair of infrastructure node types is expressed in Eq. (14). Eq. (14) is a piece-wise function that depends on the type of node hosting the different VNFs of the NS. The delay between two mobile nodes of the same cluster is accounted in $D_{M_q}(N_i, N_j)$. Delay between two nodes of the fixed infrastructure is $D_{AP,S}(N_i, N_j)$, while delay between a node of the fixed infrastructure and a mobile cluster depends on the AP that the cluster uses in that moment, which is gathered in

$\sum_{AP_k \in V_{AP}(G_I)} AP_k^q(t_u) [D_{M_q}(N_i, N^{(r_q)}) + d_{AP_k} + D_{AP,S}(AP_k, N_j)]$. Thus, the delay of a service is composed by the different delays between the nodes that host the different VNFs and the order in which they must be performed.

The overall delay of a SFC $G_s \in \mathcal{C}_{SFC}$ in time t_u is formulated in Eq. (9), where the delays between the hosts of each SFC edge are summed. The upper bound of the SFCs' total permitted delay Δ_{G_s} for the whole optimization interval is expressed in constraint Eq. (10).

3.2.3 Battery constraints

In order to place VNFs in mobile nodes it is necessary to ensure the mobile node will not run out of battery during the time interval (t_0, t_1) . This is introduced in the problem formulation, in Eq. (11), as the probability of having battery for the whole time interval considered, based on the resources used in the node. C_{N_i} is the consumed capacity of mobile node N_i , and $\mathbb{P}_{bat}(N_i, C_{N_i})$ is the probability of having battery on N_i by the end of time interval (t_0, t_1) when using C_{N_i} resources as allocated capacity. Note that Optimization problem 1 is agnostic of the used battery consumption model, as $\mathbb{P}_{bat}(N_i, C_{N_i})$ values could be derived by any battery consumption model. For example, Sec. 5 derives $\mathbb{P}_{bat}(N_i, C_{N_i})$ as a linear function between the empty $C_{N_i} = 0$ and the fully loaded states $C_{N_i} = \bar{C}_{N_i}$. To ensure the proper performance of the mobile nodes, the battery life is guaranteed in Eq. (12) by a threshold th_{bat}^s given per SFC G_s , for all nodes hosting VNFs. This threshold takes into account the battery of all the mobile nodes hosting the VNFs of the service and guarantees each of the nodes hosting a VNF of the service will have battery during the whole time interval with a probability higher than the threshold, for example a $th_{bat}^s = 0.9$.

3.2.4 Cost minimization

Finally, the problem minimizes the total cost of allocating the whole service G_s demanded and AP usages by all of the mobile clusters. Hence, the objective function is shown in Eq. (13). The VNF mapping μ and AP selection structures α are defined by the variables $x(v, N_i)$ and $AP_k^q(t_u)$ of a solution to the optimization problem. This model is not linear in some equations as the one representing the delay in Eq. (9), but each product of two variables can be easily linearized due to the fact that all the variables involved are binary variables. Thus, the linearization is performed by substituting each product of two binary variables by one extra binary variable, as expressed in Property 1.

Property 1 (Linearization of the product of two binary variables). *Let $z = x \cdot y$, where z , x and y are binary, then the product can be linearized as follows:*

$$\begin{aligned} z &\leq x, \\ z &\leq y, \\ z &\geq x + y - 1 \end{aligned}$$

Formulation 2 Bin Packing with Usage Cost [10]

Input: VNFs $V(G_S)$ as items with weight, host nodes $V_H(G_I)$ as bins with capacity

Output: VNF placement respecting only capacity constraints

$$\sum_{N_i \in V_H(G_I)} x(v, N_i) = 1 \quad \forall v \in V(G_S) \quad (15)$$

$$\sum_{v \in V(G_S)} x(v, N_i) C_v \leq \bar{C}_{N_i} \quad \forall N_i \in V_H(G_I) \quad (16)$$

$$x(v, N_i) \in \{0, 1\} \quad \forall v \in V(G_S), N_i \in V_H(G_I) \quad (17)$$

$$\min \sum_{N_i \in V_H(G_I)} C_{N_i} \cdot p_{N_i} \quad (18)$$

4 HEURISTIC

This section details the design of the heuristic which exploits the peculiarities of the system model to design an efficient and practical algorithm.

4.1 Proposed heuristic

The core idea of our heuristic algorithm is to use the fractional optimal solution of a bin packing problem of the VNFs and host nodes, which is deterministically rounded to an invalid integer solution. Next, the algorithm iteratively resolves the capacity, delay, battery and coverage constraint violations by changing the mapping location of VNFs in the initial invalid integer solution until a feasible mapping is found.

First, we introduce the bin packing problem variation with variable bin and item sizes supporting linear usage costs [10] in Formulation 2. Lemma 1 (taken from [10] and pasted down for readability) states how to construct a fractional optimal solution for this bin packing variant, relaxing the integrality constraint. The proof of Lemma 1 can be found in the original source [10].

Lemma 1 (Fractional optimal solution of Formulation 2 [10]). *Let $\{a_i\}$ be a permutation of all host infrastructure nodes $N_i \in V_H(G_I)$ in ascending order by their unit costs of computation capacity $p_{a_1} \leq p_{a_2} \leq \dots \leq p_{a_{|V_H(G_I)|}}$. Let $W_C = \sum_{v \in V(G_S)} C_v$ be the sum of all VNF capacities. Let b be the minimum number of host nodes in order $\{a_i\}$ where $\sum_{i=1}^b \bar{C}_{N_{a_i}} \geq W_C$.*

The fractional optimal solution (discarding the integrality constraint (17)) of Formulation 2 is

$$\tilde{x}(v, N_{a_i}) = \begin{cases} \frac{\bar{C}_{N_{a_i}}}{W_C} & \text{if } i < b, \\ \frac{W_C - \sum_{i=1}^{b-1} \bar{C}_{N_{a_i}}}{W_C} & \text{if } i = b, \\ 0 & \text{if } i > b; \end{cases} \quad \forall v \in V(G_S).$$

The proposed heuristic's core pseudo-code is shown in Algorithm 1. Intuitively, the heuristic reallocates VNFs that violate any constraint, and measures the goodness of the reallocation with the *improvement score* (see Algorithm 3). The higher the *improvement score*, the better the VNF reallocation. Initially, the fractional optimal solution is retrieved and rounded to initial constraint-violating VNF placement,

obeying only the locality constraints (6) as shown in lines 2-5. The cost increasing order $\{a_i\}$ of mobile and server nodes are used from Lemma 1 to involve additional hosts to the VNF placement pool, starting only from the first b cheapest hosts. In each iteration a set of violating items, respecting all constraints is calculated based on the temporary decisions stored in the current VNF placement function μ . Next, the iteration in lines 10-18 collects improvement scores for moving a VNF which is involved in any constraint violation to any currently considered host node (i.e. until index b'). Line 13 heuristically filters only the VNF relocations whose improvement score is higher than a configured improvement score limit Υ . The improvement cost is calculated by the cost difference of VNF v on the current host $\mu(v)$ and the possible new host N_{a_i} . If any allowed VNF replacement is found, update actions are taken and a current AP selection α is retrieved as shown in lines 20-22. Otherwise, the algorithm exits the improvement operations, and the next cheapest mobile or server node is included in the search by increasing b' . If a feasible solution is found after any inner iteration (see line 28), the procedure returns the current VNF placement function μ and AP selection structure α . The presented algorithm could be easily extended to continue searching for better quality solutions at the price of increased running time.

All subsequently presented subroutines take the input of Algorithm 1, but these are omitted from the pseudocodes for readability. `VIOLATINGVNFMAPPINGS` takes as input the current VNF placement function μ and returns a set of violating VNFs $\mathcal{V}$ and an information storage of the actual constraint violations $\mathcal{R}$. Based on the current VNF placement μ , the feasibility of AP selection for each robot cluster $q \in \{1 \dots Q\}$ is checked using the subroutine `CHOOSEAPS`. If the AP selection is not possible, all VNFs of the causing SFC G_s are added to $\mathcal{V}$ and the violation information is stored in constraint violation record $\mathcal{R}$.

Algorithm 2 shows the details of how the AP selection and its feasibility based on the placement function μ are derived for a given robot cluster $q \in \{1 \dots Q\}$ for all temporal subintervals. Line 4 chooses the affected SFCs G_s , which have any VNF mapped to the mobile nodes of the robot cluster q . Given the current VNF placement μ , the total delay used by the path of the whole SFC $E(G_s)$ can be calculated using the delay expression (14). Access points are chosen by discarding the ones which do not meet the coverage requirement κ_q and finding the one with minimal delay among the remaining ones:

$$AP_l = \operatorname{argmin}_{AP_k \in V_{AP}(G_I) \cap \{AP_\phi: \mathbb{P}_{AP_\phi^q}(t_u) \geq \kappa_q\}} (d_{AP_k}) \quad (19)$$

These operations are done by the function `DELAYDIS-TWITHCOVERAGEANDAPSELECTION`, which also ensures that the same AP is chosen for a given input robot cluster $q \in \{1 \dots Q\}$ in subinterval t_u , no matter which input SFC it gets. The algorithm discards the impractical option of placing the VNFs of a single SFC to distinct mobile clusters. This simplification is only applied for the delay bounded VNFs, not to the other VNFs of the network service G_S . If an access point AP_l is found for subinterval t_u with the given requirements, the selection is saved in AP selection function α , otherwise the structure is invalidated and the

Algorithm 1

Input: service graph G_S , infrastructure G_I , improvement score limit Υ , and all constraints from Sec. 3
 Output: VNF placement $\mu : V(G_S) \mapsto V_H(G_I)$ and AP selection $\alpha : \{t_u\} \times \{1 \dots Q\} \mapsto V_{AP}(G_I)$ satisfying all constraints

```

1: procedure PLACEVNFSSELECTAPS( $G_S, G_I, \Upsilon$ )
2:    $\tilde{x}(v, N_i), b, \{a_i\} \leftarrow$  fractional solution based on
   Lemma 1 for host nodes  $V_H(G_I)$  and VNFs  $V(G_S)$ 
3:   for  $v \in V(G_S)$  do ▷ Round initial solution
4:      $\mu(v) \leftarrow \operatorname{argmax}_{N_i \in V_H(G_I)} \tilde{x}(v, N_i)$  which
       obeys locality constraints (6)
5:   end for
6:   for  $b' \in \{b \dots |V_H(G_I)|\}$  do ▷ In order of  $\{a_i\}$ 
7:      $\mathcal{V}, \mathcal{R} \leftarrow \operatorname{VIOLATINGVNFMAPPINGS}(\mu)$ 
8:     while  $\mathcal{V} \neq \emptyset$  do
9:        $\mathcal{I} \leftarrow \emptyset$  ▷ Allowed improving VNF moves
10:      for  $v \in \mathcal{V}, i \in \{1 \dots b'\}$  do
11:        if  $\mu(v) \neq N_{a_i}$  and  $\mu(v) = N_{a_i}$  obeys
12:          locality constraints (6) then
13:            if  $\Upsilon \leq \operatorname{IMPROVESCORE}(\mu, v, N_{a_i})$  then
14:               $\operatorname{impr\_cost} \leftarrow C_v(p_{N_{a_i}} - p_{\mu(v)})$ 
15:               $\mathcal{I} \leftarrow \mathcal{I} \cup \{(v, N_{a_i}, \operatorname{impr\_cost})\}$ 
16:            end if
17:          end if
18:        end for
19:      if  $\mathcal{I} \neq \emptyset$  then
20:         $\mu(v) \leftarrow N_{a_i} \mid (v, N_{a_i}, \operatorname{impr\_cost}) \in \mathcal{I} \text{ and } \operatorname{impr\_cost} \text{ is minimal}$ 
21:         $\mathcal{V}, \mathcal{R} \leftarrow \operatorname{VIOLATINGVNFMAPPINGS}(\mu)$ 
22:         $\alpha \leftarrow$  retrieve AP selection from violation
          record  $\mathcal{R}$ 
23:      else break
24:      end if
25:    end while
26:    if AP selection  $\alpha$  is valid and
27:      VNF placement  $\mu$  is valid then
28:      return  $\mu, \alpha$  ▷ Solution found
29:    end if
30:  end for
31:  return  $\emptyset, \emptyset$  ▷ Solution not found
32: end procedure

```

reason is saved in $\mathcal{R}^{AP}$, as shown by the logical structure starting at line 6. In case the computation capacities of a robot cluster are not used by any VNFs of any SFC, an access point still needs to be selected for the cluster, which is done by minimizing the cost instead of the unbounded delay and similarly filtering to the coverage probability (see line 13).

Finally, the improvement score calculation is shown in Algorithm 3, which takes the current VNF placement μ and a possible relocation of VNF v to N_{a_i} as input, and outputs an integer whose higher value represents a more significant improvement. The `IMPROVESCORE` procedure uses the previously presented `VIOLATINGVNFMAPPINGS` function to evaluate how the mapping would change by the VNF mapping modification. The mapping structure μ with less violating constraints is considered better, as shown in lines modifying the improvement score y . In case of capacity

Algorithm 2

Input: Current VNF placement μ , current (possibly incomplete or invalid) AP selection α , robot cluster index q
 Output: Extended and/or invalidated AP selection α , AP selection violation record $\mathcal{R}^{AP}$

```

1: procedure CHOOSEAPS( $\mu, \alpha, q$ )
2:   for  $t_u \in (t_0, t_1)$  do
3:     if  $\exists G_s \in \mathcal{C}_{SFC}, \exists v \in V(G_s),$ 
4:       where  $\mu(v) \in V_{RC_q}(G_I)$  then
5:        $d^{G_s}, AP_l \leftarrow \operatorname{DELAYDISTWITHCOVERAGE-}$ 
         $\operatorname{ANDAPSELECTION}(E(G_s), \mu, t_u, q, \kappa_q)$ 
6:       if  $d^{G_s} \leq \Delta_{G_s}$  and  $\exists AP_l$  then
7:         Let  $\alpha(t_u, q) = AP_l$  ▷ Same AP for all SFCs
8:       else
9:         Let  $\alpha(t_u, q) = \lceil$ 
10:        Add result  $d^{G_s}$  and SFC  $G_s$  to  $\mathcal{R}^{AP}$ 
11:      end if
12:    else
13:      Let  $\alpha(t_u, q) = AP_l$ , where  $AP_l \in V_{AP}(G_I)$ 
        and obeys coverage constraint (8) and  $p_{AP_l}$ 
        is minimal
14:    end if
15:  end for
16:  return  $\alpha, \mathcal{R}^{AP}$ 
17: end procedure

```

constraints, total improvement score y would decrease, keep unchanged or increase if the number of hosts with more than their max capacity allocated would increase, stay or decrease by the VNF movement, respectively (see line 5). A similar score modification is done for each SFC, using the change in the number of temporal subintervals t_u where the coverage or delay constraints are violated as shown by the iteration starting at line 6. In the case of the battery constraints, the number of VNFs mapped to mobile nodes with violated battery thresholds are used.

4.2 Complexity analysis

A brief analysis on the heuristic's complexity and its termination is presented in Theorem 1 and its corresponding proof.

Theorem 1 (Complexity of heuristic). *The overall complexity of the heuristic with positive improvement score limit $\Upsilon > 0$ is:*

$$\mathcal{O}\left(|V(G_S)|^4 |V(G_I)|^3 |\mathcal{C}_{SFC}| QT\right) \quad (20)$$

where Q and T are the number of clusters and the number of subintervals t_u in the optimization time frame (t_0, t_1) , respectively.

Proof. Looking at Algorithm 1, the fractional solution construction and its rounding are dominated by the iteration starting at line 6, which is executed at most $|V(G_I)|$ times. Assuming a positive improvement score limit Υ , the violating VNFs set $\mathcal{V} = \mathcal{O}(|V(G_S)|)$ decreases at least by one element in each iteration of the *while* cycle. At most every iteration runs `VIOLATINGVNFMAPPINGS`. Filtering for the allowed VNF movements in line 12 is done at most $\mathcal{O}(|V(G_S)||V(G_I)|)$ times, and in worst case for each of them we execute a `IMPROVESCORE` subroutine

Algorithm 3

Input: Current VNF placement μ , movement of VNF v to host N_{a_i}

Output: Integer in interval $[-|\mathcal{C}_{SFC}| - 2, |\mathcal{C}_{SFC}| + 2]$, the improvement score of the VNF movement

```

1: procedure IMPROVESCORE( $\mu, v, N_{a_i}$ )
2:    $y \leftarrow 0 \triangleright$  Init. improvement score of moving  $v$  to  $N_{a_i}$ 
3:    $\mathcal{V}, \mathcal{R} \leftarrow \text{VIOLATINGVNFMAPINGS}(\mu)$ 
4:    $\mathcal{V}', \mathcal{R}' \leftarrow \text{VIOLATINGVNFMAPINGS}(\mu \mid \mu(v) = N_{a_i})$ 
5:    $y \leftarrow y - 1 / + 0 / + 1$  if number of hosts  $N_i$  with
     violated constraint (4) increases/stays/decreases in
      $\mathcal{R}'$  compared to  $\mathcal{R}$ 
6:   for  $G_s \in \mathcal{C}_{SFC}$  do
7:      $y \leftarrow y - 1 / + 0 / + 1$  if number of subintervals  $t_u$ 
       with any invalid mappings (i.e. where  $\exists t_u, q :$ 
        $\alpha(t_u, q) = \lceil \cdot \rceil$ ) increases/stays/decreases in  $\mathcal{R}'$ 
       compared to  $\mathcal{R}$ 
8:   end for
9:    $y \leftarrow y - 1 / + 0 / + 1$  if number of VNFs  $v$  which are
     mapped to any mobile node  $V_M(G_I)$  with violated
     battery constraint (12) increases/stays/decreases in
      $\mathcal{V}'$  compared to  $\mathcal{V}$ 
10:  return  $y$ 
11: end procedure

```

call. These observations make Algorithm 1's complexity to be $\mathcal{O}(|V(G_S)||V(G_I)|\{\mathcal{O}(|V(G_S)||V(G_I)|\text{IMPROVESCORE} + \text{VIOLATINGVNFMAPINGS})\})$. The VIOLATINGVNFMAPINGS's complexity is dominated by $Q\mathcal{O}(\text{CHOOSEAPS})$, because the other constraints can be checked in $\mathcal{O}(|V(G_I)||V(G_S)|)$ time. Access point filtering for sufficient coverage in a longest SFC can be done in $\mathcal{O}(|V(G_I)||V(G_S)|)$ time, which is done for all SFCs $|\mathcal{C}_{SFC}|$, all SFC edges $\mathcal{O}(|V(G_S)|)$ for all time subintervals T . Which gives $\mathcal{O}(\text{VIOLATINGVNFMAPINGS}) = \mathcal{O}(QT|V(G_S)|^2|V(G_I)||\mathcal{C}_{SFC}|)$. Similarly, IMPROVESCORE is dominated by VIOLATINGVNFMAPINGS's complexity. Finally, a Floyd-Warshall algorithm is used to pre-calculate the all the delay matrices $D_{AP,S}$ and D_{M_q} with complexity $\mathcal{O}(|V(G_I)|^3)$, which is dominated by the previous operations. Substituting and ordering the $\mathcal{O}(\cdot)$ notations, the statement follows. $\square$

5 EVALUATION AND RESULTS

This section compares the performance of Sec. 4 heuristic, with the optimal solution of Sec. 3 formulation from various aspects. As integer programs are generally impractical due to the hardness of the problem, our heuristic is extensively evaluated to demonstrate its applicability. The heuristic solutions are compared to the optimal solution obtained with Gurobi which finds a solution within a gap optimality of 3%. Such comparison is done for the mobile robotics use case of Sec. 2, where scalability is a critical issue due to the size of the infrastructure and service graphs.

Additionally, this section compares Sec. 4 heuristic against "Follow Me Chain" (FMC) [11], a heuristic that tackles mobility by triggering VNF migrations upon AP handovers, but does not consider battery constraints. Our

Fig. 2: A service graph generated with a series-parallel graph. This instance contains x8 VNFs bounded to mobile nodes, and is used in the battery experiment, see Fig. 4

implementation of FMC (i) replaces [11, Algorithm 1] VNF-based Breadth-First Search (BFS) with a virtual-link-based BFS, so as to ensure the mapping of every virtual link; (ii) uses a k -shortest paths in [11, Algorithm 2:line 1] to avoid getting stuck in the search of all paths between two nodes¹; (iii) considers mobile compute nodes as well as edge servers; and (iv) can map service graphs with unconnected components.

5.1 Experiment setup

The presented evaluation scenario scales up the mobile robotics use case of Valencia city haven. A realistic 5G network infrastructure topology is considered with multiple types of wireless access points, while the service graph instances are random graphs. Many parameters of the experiment setting are examined during the presented simulations, varying the size of the input, SFC delay requirements, coverage probabilities and battery thresholds.

In order to generalize the service graphs and gain confidence in our simulations, series-parallel graphs are used to generate the network service topology G_S . Fig. 2 shows an example of such graph. This graph class covers the structure of many data streaming applications, such as map-reduce topologies, and have been used in other realistic, industrial case studies for fog application allocation [12]. The round-trip time experienced by every robot running a VNF must stay below the delay restriction, therefore, SFCs correspond to loops starting and ending in mobile node VNFs, and every SFC must satisfy the delay restriction. Among all the VNFs of the SFC, some of them are forced to run on top of the mobile robots' hardware $V_M(G_I)$ (denoted as *Mobile node VNF* in Fig. 2), and the rest can run on top of any server $V_S(G_I)$ or robot $V_M(G_I)$. It is up to the heuristic and the optimization formulation, to decide where to deploy them.

Every experiment uses the 5G infrastructure characterization of [7] and [6], which considers Ultra Reliable Low Latency Communications. Table 2 shows every infrastructure element considered in the experiments, and Fig. 1 illustrates the interconnection of the network infrastructure. Each M1 switch is located in the access ring of the network, and it gathers the traffic of up to x6 LTE or NR RUs.

¹ FMC builds a full-mesh servers' graph, and even the proposed range-based Depth-First Search (DFS) incurs into a $\mathcal{O}(|V(G_i)|!)$ search space

TABLE 2: Infrastructure used for experimentation

#	Element	Characteristics
x2	LTE RU [16]	8km radio coverage, 5ms one way delay [17], 5.5 cost units OPEX [18]
x36	NR RU [19]	700m LoS coverage [20], 1ms one way delay, 11 cost units OPEX
x10	robots	x2 CPUs, 15.27 cost units/CPU [21]
x6	edge server	x12 CPUs, 5.83 cost units/CPU [21]
x2	cloud rack	200 CPUs, 2.46 cost units/CPU [21]
x8	M1 switch	x8 dedicated CPUs
x6	M2 switch	x238 dedicated CPUs
x2	access rings	fiber ring connection, ≤ 6 M1 switches
x1	aggregation ring	fiber ring connection, x6 M2 switches

Access rings have x6 M1 switches and x1 M2 switch, all of them interconnected in a ring fashion. Every M2 switch belongs to x4 access rings, and it steers the traffic up to the aggregation ring, where it is connected in a ring fashion with another x5 M2 switches. Experiments consider that edge and cloud servers are reachable using M1 and M2 switches, respectively.

Each point in the operation area of the robot cluster is covered at least by one LTE RU and at least by one NR RU. The coverage probabilities for each time instance are derived by a function which maps the distance of the RU and the cluster to the coverage probability. The probability slightly decreases until the end of the RU coverage area, and steeply drops to 0 at 120% of the RU reach. If a NR RU and the mobile cluster are not in LoS, the coverage probability is 0, independent of their distance. To achieve e2e delays demanded by the mobile robotics use case (between 1ms and 100ms), the experiment infrastructure assumes that aggregation and access ring switches introduce packet processing delays between 1ms and 10ms, under the same characterisation as performed in [7].

To derive this section's results, a network infrastructure with just one cluster of robots has been generated with the 5GEN R package [13]. Then, a Python script generates series-parallel NS graphs G_S from which loop SFCs are selected. Robot cluster paths are encoded by coordinates which are used to calculate RUs coverage probabilities as robots move along the path. Next, the Python script runs Sec. 4 heuristic to decide each VNF mapping on top of the infrastructure graph. Sec. 3.2 formulation is encoded in AMPL [14], and the Python script invokes Gurobi 8.1 solver [15] through the *amplpy* API to obtain the optimal mapping. All the experiments have been executed on two identical VMs with x4 vCPUs, 32GB of memory, and 132GB of disk.

5.2 Simulation results

This section presents the results of the extensive simulations performed with Algorithm 3, AMPL solver, and the state of the art FMC solution; denoted as *impr- Υ* , *AMPL*, and *FMC*; respectively. The details of the simulation parameters are shown in Table 3 for each of the experiments. The cluster paths in the Valencia haven are represented by their source and target locations.

All evaluation figures present boxplots, where the middle line shows the median (a.k.a. second quartile) of the dataset, while the body of the boxplots show the first and third quartiles (a.k.a. the medians of the first half and

the second half of the dataset separated by its median). The whiskers of the boxplots represent the datum which deviates from the boxplot body at most by 1.5 times the inter-quartile range, while outliers are individually plotted by circles which fall beyond the whiskers.

An input VNF placement problem with all previously presented constraints is deemed feasible, if the AMPL implementation finds a valid solution that respects all constraints in 30 minutes (measured in wall-clock time). In case of the heuristic, the timeout is reduced to 20 minutes. All experiments were executed with 3% optimality gap for AMPL, various improvement score limit values for the heuristic, and $k = 10$ for FMC.

First of all, the scalability of the algorithms are compared depending on the number of VNFs to be placed; results in terms of cost and runtime are shown in Fig. 3a and Fig. 3d, respectively. The time-bound feasibility is shown on top of the figures for each randomized scenario repetition corresponding to the dependant value on the horizontal axis. The scalability experiment is repeated multiple times for each input size, varying the distribution of VNF capacity requirements, the service graph's concrete topology, and the selection of the VNFs bound to the mobile cluster (see Table 3). The scenario parameters allow a solution to be found in any randomized generation, due to loose SFC delay, coverage and battery probability thresholds; though the 30mins time limit may not be enough in all cases. Fig. 3a shows the time-bound feasibility ratio calculated on the randomized repetitions. A steep drop of feasibility of the AMPL implementation occurs at the VNF count of 60, which is due to reaching the computation timeout in each case. The reason behind the timeouts is the exponential runtime of the AMPL solution, which is shown by Fig. 3d in logarithmic time scale. On the contrary, both heuristics find feasible solutions in every possible setup. In terms of cost, our heuristic with $\Upsilon = 1$ outperforms FMC, with the former staying between 15% and 30% away of the optimal costs, and the latter increasing the cost gap with respect to our solution as the number of NFs bound to mobile nodes grows. Furthermore, our heuristic always find solutions below 100ms for all tests, whilst FMC takes around 10s.

Second, the effect of the coverage probability threshold κ_q is studied. Fig. 3b shows how the cost varies by increasing the threshold, i.e. making the AP selection more strict. As the coverage probability requirement increases, deployment costs become more expensive, because the solutions impose the selection of the closer and more expensive NR antennas, rather than the cheap LTE antennas. Fig. 3b depicts as well the feasibility, and shows that for $\kappa_q = 0.99$ all scenarios are infeasible, because there exists at least one subinterval in which the cluster is not covered by any antennas with such high probability. Regarding the impact of the improvement score Υ , Fig. 3b and Fig. 3e show that $\Upsilon = 2$ (*impr-2* time series) finds cheaper solutions faster. This is due to the heuristic design, which goes faster by shrinking the solution space and considering only VNF relocations with higher improvement score. The heuristic finds cheaper deployments faster, because they require less steps to make the rounded fractional solution feasible. Additionally, Fig. 3b shows that FMC cannot find feasible solutions with the studied coverage thresholds $\kappa_q \geq 0.9$, since one or more

TABLE 3: Experiment parameters

Parameter name/explanation	Value/range			
Experiment name	Scalability	Coverage	Delay	Battery
Robot cluster path, see Fig. 1	$S \rightarrow D1$	$S \rightarrow D2$	$S \rightarrow D2$	$S \rightarrow D1$
Path total distance [meters]	678	488	488	678
Time interval count $t_u \in (t_0, t_1)$	24	24	24	24
Unloaded battery probability $\mathbb{P}_{bat}(N_i, 0), \forall N_i \in V_M(G_I)$	99%	99%	99%	99%
Full loaded battery probability $\mathbb{P}_{bat}(N_i, \bar{C}_{N_i}), \forall N_i \in V_M(G_I)$	50%	80%	80%	50%
Battery probability threshold th_{bat}^s	40%	70%	70%	72% & 75%
Infrastructure delay sample count	1	4	4	1
SFC delay [ms] Δ_{G_S}	1000	5	Varies	1000
Randomized VNF vCPU requirement $C_v \forall v \in V(G_S)$	$0.5 \times \{0, \dots, 4\}$	$0.5 \times \{0, \dots, 4\}$	$0.5 \times \{0, \dots, 4\}$	$0.25 \times \{0, \dots, 4\}$
VNF count $ V(G_S) $	Varies	10	10	26
VNF count bound to robots	6	4	1	Varies
Coverage probability threshold κ_q	94%	Varies	94%	70%
Scenario repetition with different randomization seed	14	24	20	14

Fig. 3: Results of scalability, coverage probability and SFC delay experiments

migrations failed during the experienced handovers.

Next, the results of simulations varying the SFC delay are shown in Fig. 3c and Fig. 3f. FMC cannot find feasible solutions for 3ms scenarios, as it is designed to try to map one VNF per compute node, and therefore, its mappings have to traverse more network links. The heuristic *impr-1* struggles with finding feasible solutions in the allocated time for the 3ms scenarios, while AMPL manages to prove the existence of valid solutions as shown by the feasibility percentages of Fig. 3c. This could be easily addressed by introducing a search space pruning step in addition to the locality constraints. In the 3ms scenarios the usage of the cheap and high capacity cloud nodes is not an option because their RTTs from all APs are above this value. Excluding these compute nodes from the allocation options for

the VNFs contained in the strict SFCs would dramatically decrease the running time and thus increase the time-bound feasibility of the heuristic. Although, additional pruning steps decrease solution quality in the cases of more permissive delay requirements. On the other hand, the heuristic greatly outperforms the optimal solution search in the 10-15ms scenarios, where the AMPL algorithm fails to find any feasible solution before the 30 minutes timeout. This is due to the growth of the search space as the delay restriction is relaxed. Additionally, *impr-1* finds cheaper deployments than FMC, since the latter tries to use one compute node per VNF, and does not account for cloud nodes by design. Note that cloud nodes are cheap and strong candidates used by *impr-1* when the delay requirement relaxes (see the SFC delay case of 1000ms). Another interesting aspect of the

Fig. 4: Impact of battery probability threshold on cost and feasibility

solutions is the number of required handovers needed for the whole optimization time interval. A lower handover count requires less management operations and results in a more stable service. Handover comparison between the cost-optimal and the heuristic solutions are shown in Fig. 3f. The heuristic outperforms the optimal solution, which is especially relevant when the scenario could be solved by a few handovers as shown by the 10ms experiment scenarios with 100% heuristic feasibility. The AMPL algorithm could be modified to minimize the number of handovers, but it would further worsen its scalability, while the heuristic performs well by design. Furthermore, *impr-1* required less handovers than FMC in all the simulated scenarios.

Last, the results of the conducted experiments to examine the battery threshold parameter's effects are shown in Fig. 4. The figure depicts cost values for both algorithms in cases of 72% and 75% battery alive probability requirements, as the number of VNFs to be placed on the mobile cluster increases. Note that FMC is agnostic of battery constraints and it reports the same solution, no matter the imposed battery alive probability. However, the feasibility of the FMC solution is depicted for battery alive cases. These scenarios challenge constraint (12), discovering the critical battery threshold to be around 72%-75%. In the 72% case the scenarios are vastly feasible with a slight decrease as the VNF bound to mobile nodes increase. The heuristic finds close to optimal allocations in almost all scenarios, except in the extreme case of much freedom. In the more strict case of 75%, besides the no location-bound VNF experiment which is essentially the same as the 72% case, the heuristic always finds all optimal solutions where it exists. Last of all, Fig. 4 shows that FMC only finds solutions in the 7% of the simulations with more than 8 VNFs bound to mobile nodes. Indeed, it reports same feasibility ratios for both 72%, 75% battery thresholds, as it could only find deployments with $\mathbb{P}_{bat}(N_i, C_{N_i}) \geq 0.75$. As in previous results, FMC reports higher deployment costs because it tries to map each VNF to a different compute node.

The implementation of the algorithms, the simulation framework and all presented scenarios with raw data are available for further usage or result reproduction².

6 RELATED WORK

Due to the widespread of virtualization technologies, the problem of allocating VNFs on top of physical resources has been of interest in recent years. In most of the existing research the allocation of VNFs is envisioned as an optimization problem, that is generally $\mathcal{NP}$ -hard [22].

A common technique is to solve the VNF allocation problem as a variation of the bin packing problem, taking the VNFs as items, and the bins as servers. Particularly, the first steps of this paper's proposed heuristic are build upon the basis defined in the algorithm of [10], which minimizes a data center energy consumption using a generalized bin packing problem. Works as [23] solve the VNF allocation using the variable size bin packing problem [24], which provides an efficient solution to minimize both response time, and resource utilization. Other research projects have studied different and relevant generalizations for variable sized bin-dependent costs [25]. A recent survey categorizes bin packing problem generalizations which might be relevant to VNF placement solutions [26]. In general, algorithms for bin packing problems do not consider delays on the sequence of items, nor any topological constraint among the bins, so using their results for VNF placement problems is not trivial; our heuristic builds on such results.

Solutions of the VNF allocation problem must reshape with the new 5G networks, which bring computational capabilities closer to the user thanks to MEC [27], and fog computing [28]. Indeed, servers are way closer to antennas, or even co-located with them in the edge, and IoT devices are becoming part of a dense network. Thus, 5G comes with the urge of a more dense radio coverage, and the possibility of sharing public/private network infrastructure [2] [29] can help to achieve it. Orchestration in the edge of 5G has motivated solutions [11] that benefit from edge servers to asses the mapping and migration of VNF resources upon users' mobility. Additionally, edge computing has popped up the quest of deploying NSs with very strict latency requirements, and recent research as [30], [31], [32], [33], and [34] study solutions about how to allocate VNFs to meet low latency requirements. [30] uses a genetic algorithm to obtain a fixed allocation that minimizes/maximizes latency/availability, [34] provides a stopping theory solution that migrates the allocated VNFs as time passes, such that latency restrictions are not violated. [31] formulates an optimization problem to allocate VNFs demanded by end-users attached to antennas, so as to maximize/minimize resources re/usage, by imposing latency constraints. [32] proposes a deep learning agent that assigns VNFs to servers maximizing the requests' throughput, while they meet latency constraints. [33] presents a solution that maximizes the throughput of services in 5G slices, while meeting latency requirements of each slice. The solution idea relies on preventing the performance interference caused by co-locating multiple VNFs in the same server.

2. GitHub link will be added to the camera ready version.

There is some research that focus on VNF allocation in fog environments. [35] presents an allocation model accounting for the computational overhead of fog devices, based on the assigned workload; and [36] studies how to satisfy e2e delay by reducing the distance of the deployed service, to the users consuming the service (envisioned as traffic generators). Another approach to allocate VNFs is to deploy them jointly using cloud and fog devices, as [37] does. In that work, service providers derive a wireless and resource sharing model of fog devices, and the allocation is done using a student project allocation algorithm. There are other results [38] related to low energy IoT devices, that study the trade-off between the energy requirement for computation, and transmitting data, as a computation task outsourcing pipeline is proposed.

Although the literature already provides solutions to perform the VNF allocation on edge and fog scenarios, this paper contributes to the state-of-the-art by tackling all at once the (i) radio coverage; (ii) battery consumption; and (iii) e2e delay restrictions present in 5G use cases with mobile compute nodes.

7 CONCLUSIONS

This paper has analyzed the notoriously hard problem of VNF placement in a realistic use case based scenario: mobile robotics for warehousing solution in the Valencia city haven, where an NPN deployment in a public network is assumed. In this scenario, mobile compute nodes act as an extension of the cloud and edge computing infrastructure, which triggers the need for VNF placement solutions with strict delay bounds and reliability constraints, while taking into account radio coverage, mobility and battery conditions.

The paper has introduced a system model and a mathematical formulation of the problem, to then propose an efficient heuristic building on the fractional optimal solution of a bin packing variant. The heuristic has been extensively evaluated via simulations in terms of scalability and the strictness of constraints which are relevant to the use case. Results show that the proposed heuristic outperforms a state of the art mobility-aware algorithm, and achieves close to optimal deployments' in terms of cost, while improving the convergence speed to the solution (therefore the number of time-feasible solutions is increased) and minimizing the number of required handovers. To the best of our knowledge, our solution is the first to tackle the VNF placement problem simultaneously respecting battery, coverage and delay constraints over a mobile and volatile 5G infrastructure.

ACKNOWLEDGEMENTS

This work has been partially funded by the EC H2020 5G-DIVE Project (grant no.: 859881) and H2020-ICT-2018-3 EC H2020 5G-GROWTH Project (grant no.: H2020-ICT-2018-3 856709).

REFERENCES

- [1] 3GPP, "Study on Scenarios and Requirements for Next Generation Access Technologies," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 38.913, 08 2017, version 14.3.0.
- [2] 5GACIA, "5G Non-Public Networks for Industrial Scenarios," 5G Alliance for Connected Industries and Automation, White Paper, 7 2019.
- [3] 3GPP, "Study on Communication for Automation in Vertical Domains," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 22.804, 07 2020, version 16.3.0.
- [4] 5G-CORAL, "5G-CORAL initial system design, use cases, and requirements," Tech. Rep. deliverable D1.1, 2018.
- [5] B. Nogales, V. Sanchez-Aguero, I. Vidal, and F. Valera, "Adaptable and automated small uav deployments via virtualization," *Sensors*, vol. 18, no. 12, p. 4116, 2018.
- [6] ITU-T, "Consideration on 5G transport network reference architecture and bandwidth requirements," International Telecommunication Union - Telecommunication Standardization Sector (ITU-T), Study Group 15 Contribution 0462, 2 2018.
- [7] L. Cominardi, L. M. Contreras, C. J. Bernados, and I. Berberana, "Understanding QoS Applicability in 5G Transport Networks," in *2018 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, June 2018, pp. 1–5.
- [8] V. Suryaprakash, J. Möller, and G. Fettweis, "On the modeling and analysis of heterogeneous radio access networks using a poisson cluster process," *IEEE Transactions on Wireless Communications*, vol. 14, no. 2, pp. 1035–1047, 2015.
- [9] V. Sanchez-Aguero, F. Valera, I. Vidal, C. Tipantuña, and X. Hesselbach, "Energy-Aware Management in Multi-UAV Deployments: Modelling and Strategies," *Sensors*, vol. 20, no. 10, p. 2791, 2020.
- [10] H. Cambazard, D. Mehta, B. O'Sullivan, and H. Simonis, "Bin Packing with Linear Usage Costs - An Application to Energy Management in Data Centres," *Principles and Practice of Constraint Programming - 19th International Conference*, p. ?, 2013, best Application Paper Award. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-00858159>
- [11] Y.-T. Chen and W. Liao, "Mobility-aware service function chaining in 5g wireless networks with mobile edge computing," in *ICC 2019-2019 IEEE International Conference on Communications (ICC)*. IEEE, 2019, pp. 1–6.
- [12] M. Suter, R. Eidenbenz, Y.-A. Pignolet, and A. Singla, "Fog application allocation for automation systems," in *2019 IEEE International Conference on Fog Computing (ICFC)*. IEEE, 2019, pp. 97–106.
- [13] C. J. B. A. M. Jorge Martín-Pérez, Luca Cominardi, "5gen: A tool to generate 5g infrastructure graphs," in *2019 IEEE Conference on Standards for Communications and Networking (CSCN)*, 2019.
- [14] R. Fourer, D. M. Gay, and B. W. Kernighan, "AMPL: a modeling language for mathematical programming," 1993.
- [15] I. Gurobi Optimization, "Gurobi optimizer reference manual," URL <http://www.gurobi.com>, 2015.
- [16] 3GPP, "Overview of 3GPP Release 8," 3rd Generation Partnership Project (3GPP), Technical Specification (TS), 09 2016.
- [17] R. Stuhlfauth, "UMTS Long Term Evolution(LTE)Reiner Stuhlfauth," 11 2012.
- [18] V. Nikolicj and T. Janevski, "A cost modeling of high-capacity lte-advanced and ieee 802.11 ac based heterogeneous networks, deployed in the 700 mhz, 2.6 ghz and 5 ghz bands," *Procedia Computer Science*, vol. 40, pp. 49–56, 2014.
- [19] N. Patriciello, S. Lagen, L. Giupponi, and B. Bojovic, "5g new radio numerologies and their impact on the end-to-end latency," in *2018 IEEE 23rd International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, Sep. 2018, pp. 1–6.
- [20] B. Halvarsson, A. Simonsson, A. Elgcróna, R. Chana, P. Machado, and H. Asplund, "5g nr testbed 3.5 ghz coverage results," in *2018 IEEE 87th Vehicular Technology Conference (VTC Spring)*, June 2018, pp. 1–5.
- [21] 5G-TRANSFORMER, "5G-TRANSFORMER final system design and Techno-Economic Analysis," Tech. Rep. deliverable D1.4, 2019.
- [22] M. Rost and S. Schmid, "Charting the complexity landscape of virtual network embeddings," in *2018 IFIP Networking Conference (IFIP Networking) and Workshops*, May 2018, pp. 1–9.
- [23] Q. Zhang, Y. Xiao, F. Liu, J. C. Lui, J. Guo, and T. Wang, "Joint optimization of chain placement and request scheduling for network function virtualization," in *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2017, pp. 731–741.
- [24] D. K. Friesen and M. A. Langston, "Variable sized bin packing," *SIAM journal on computing*, vol. 15, no. 1, pp. 222–230, 1986.

- [25] M. M. Baldi, D. Manerba, G. Perboli, and R. Tadei, "A Generalized Bin Packing Problem for parcel delivery in last-mile logistics," *European Journal of Operational Research*, vol. 274, no. 3, pp. 990–999, 2019. [Online]. Available: <https://ideas.repec.org/a/eee/ejores/v274y2019i3p990-999.html>
- [26] E. G. Coffman, J. Csirik, G. Galambos, S. Martello, and D. Vigo, "Bin packing approximation algorithms: Survey and classification," 2013.
- [27] Y. F. P. K. e. a. Sami Kekki, Walter Featherstone, "MEC in 5G networks," European Telecommunications Standards Institute (ETSI), White paper 28, 6 2018.
- [28] Z. Wen, R. Yang, P. Garraghan, T. Lin, J. Xu, and M. Rovatsos, "Fog orchestration for internet of things services," *IEEE Internet Computing*, vol. 21, no. 2, pp. 16–24, 2017.
- [29] 3GPP, "Service requirements for the 5G system; Stage 1," 3rd Generation Partnership Project (3GPP), Technical Specification (TS), 06 2018, version 15.5.0.
- [30] L. Yala, P. A. Frangoudis, and A. Ksentini, "Latency and availability driven vnf placement in a mec-nfv environment," in *2018 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2018, pp. 1–7.
- [31] P. Jin, X. Fei, Q. Zhang, F. Liu, and B. Li, "Latency-aware vnf chain deployment with efficient resource reuse at network edge," in *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*. IEEE, 2020, pp. 267–276.
- [32] Y. Xiao, Q. Zhang, F. Liu, J. Wang, M. Zhao, Z. Zhang, and J. Zhang, "NFVdeep: Adaptive online service function chain deployment with deep reinforcement learning," in *Proceedings of the International Symposium on Quality of Service*, 2019, pp. 1–10.
- [33] Q. Zhang, F. Liu, and C. Zeng, "Adaptive interference-aware VNF placement for service-customized 5G network slices," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 2449–2457.
- [34] R. Cziva, C. Anagnostopoulos, and D. P. Pezaros, "Dynamic, latency-optimal vnf placement at the network edge," in *IEEE infocom 2018-ieee conference on computer communications*. IEEE, 2018, pp. 693–701.
- [35] K. Intharawijitr, K. Iida, and H. Koga, "Analysis of fog model considering computing and communication latency in 5g cellular networks," in *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*. IEEE, 2016, pp. 1–4.
- [36] A. Santoyo-González and C. Cervelló-Pastor, "Latency-aware cost optimization of the service infrastructure placement in 5g networks," *Journal of Network and Computer Applications*, vol. 114, pp. 29–37, 2018.
- [37] Y. Gu, Z. Chang, M. Pan, L. Song, and Z. Han, "Joint radio and computational resource allocation in iot fog computing," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 8, pp. 7475–7484, 2018.
- [38] F. Samie, V. Tsoutsouras, L. Bauer, S. Xydis, D. Soudris, and J. Henkel, "Computation offloading and resource allocation for low-power iot edge devices," in *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, Dec 2016, pp. 7–12.

Balázs Németh is an Industrial Ph.D. student at Budapest University of Technology and Economics (BME) in cooperation with Ericsson Research. He obtained his M.Sc. degree at BME as a Computer Science Engineer in information communication specialization (2016). He has been working on orchestration algorithms for the H2020 5G-PPP 5G Exchange (5GEx) project. Currently, he is pursuing his Ph.D. degree in network softwareization with special focus on orchestration algorithms and next generation network

models.

Nuria Molner obtained her BSc in Mathematics from University of Valencia (Spain) in 2015 and her MSc in Telematics Engineering in University Carlos III of Madrid (Spain) in 2016. Currently, she is studying her PhD in Telematics Engineering from University Carlos III of Madrid, in the topic of optimization of 5G networks, while developing her research work at IMDEA Networks Institute and participating in European H2020 5G projects.

Jorge Martín Pérez is a telematics engineering Ph.D. student at Universidad Carlos III de Madrid. He obtained a BSc in mathematics in 2016, and a BSc in computer science in 2016, both at Universidad Autónoma de Madrid. Latter on 2017 he obtained a M.Sc. in telematics engineering in Universidad Carlos III de Madrid. His research focuses in optimal resource allocation in 5G networks, and since 2016 he participates in EU funded research projects.

pers in international journals and conferences. He has participated in several EU funded projects, being the project coordinator of 5G-TRANSFORMER and 5Growth.

Antonio de la Oliva received his telecommunications engineering degree in 2004 and his Ph.D. in 2008 from the Universidad Carlos III Madrid (UC3M), Spain, where he has been an associate professor since then. He is an active contributor to IEEE 802 where he has served as Vice-Chair of IEEE 802.21b and Technical Editor of IEEE 802.21d. He has also served as a Editor of several journals and magazines. He has published more than 50 papers on different networking areas. He has participated in several

Eu funded project and he is currently coordinating the EC/TW joint project 5G-DIVE.

Balázs Sonkoly is an associate professor at Budapest University of Technology and Economics (BME) and he is the head of MTA-BME Network Softwareization Research Group. He received his Ph.D. (2010) and M.Sc. (2002) degrees in Computer Science from BME. He has participated in several EU projects (FP7 OpenLab, FP7 UNIFY, H2020 5G Exchange) and national projects. He was the demo co-chair of ACM SIGCOMM 2018, EWSDN'15, IEEE HPSR'15. His current research activity focuses on cloud / edge / fog

computing, NFV, SDN, and 5G.