

Received February 26, 2017, accepted March 14, 2017, date of publication March 21, 2017, date of current version April 24, 2017.

Digital Object Identifier 10.1109/ACCESS.2017.2685437

Migration Energy Aware Reconfigurations of Virtual Network Function Instances in NFV Architectures

VINCENZO ERAMO¹, MOSTAFA AMMAR², (Fellow, IEEE),
AND FRANCESCO GIACINTO LAVACCA¹

¹DIET, Sapienza University of Rome, 00184 Rome, Italy

²College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280 USA

Corresponding author: V. Eramo (vincenzo.erao@uniroma1.it)

ABSTRACT Network function virtualization (NFV) is a new network architecture framework that implements network functions in software running on a pool of shared commodity servers. NFV can provide the infrastructure flexibility and agility needed to successfully compete in today's evolving communications landscape. Any service is represented by a service function chain (SFC) that is a set of VNFs to be executed according to a given order. The running of VNFs needs the instantiation of VNF instances (VNFI) that are software modules executed on virtual machines. This paper deals with the migration problem of the VNFI needed in the low traffic periods to turn OFF servers and consequently to save energy consumption. Though the consolidation allows for energy saving, it has also negative effects as the quality of service degradation or the energy consumption needed for moving the memories associated to the VNFI to be migrated. We focus on cold migration in which virtual machines are redundant and suspended before performing migration. We propose a migration policy that determines when and where to migrate VNFI in response to changes to SFC request intensity. The objective is to minimize the total energy consumption given by the sum of the consolidation and migration energies. We formulate the energy aware VNFI migration problem and after proving that it is NP-hard, we propose a heuristic based on the Viterbi algorithm able to determine the migration policy with low computational complexity. The results obtained by the proposed heuristic show how the introduced policy allows for a reduction of the migration energy and consequently lower total energy consumption with respect to the traditional policies. The energy saving can be on the order of 40% with respect to a policy in which migration is not performed.

INDEX TERMS Network function virtualization, migration policy, power consumption, Viterbi algorithm.

I. INTRODUCTION

Network Function Virtualization (NFV) is emerging as a new network architecture that uses standard IT virtualization techniques to consolidate many network equipment types onto industry standard high volume servers [1]. The service functions are virtualized and executed on Virtual Network Function Instances (VNFI), that are Virtual Machines running in Commercial-Off-the-Shelf (COTS) servers. The virtualization makes service deployment more flexible and scalable. The efficiency of a network, however, closely depends on the placement of the VNFI as well as the routing of service function chains (SFCs), characterizing a set of service functions to be executed for a packet flow. For this reason efficient algorithms have to be introduced to determine where

to instantiate the VNFI and to route the SFCs by choosing the network paths and the VNFI involved. The algorithms have to take into account the limited resources of the network links and servers and pursue objectives of load balancing, energy saving and recovery from failure [2]. The task of placing SFC is closely related to virtual network embeddings [3] and may, therefore, be formulated as an optimization problem, with a particular objective. The main proposed solutions are discussed in [2]. For instance Moens and Turck [4] formulate the SFC placement problem as a linear optimization problem that has the objective to minimize the number of active servers. The optimization problem and an heuristic have been also introduced in [5] where the processing resource in terms of the number of cores assigned to each VNFI is also evaluated.

In this paper we consider situations when VNFI placement needs to be changed as traffic demands change over time. Such change of placement, called VNFI *migration*, is desirable for several reasons. Function instances may be migrated to new locations to adapt to changing traffic patterns. Also VNFI consolidation in fewer servers can provide significant energy savings by consolidating the VNFIs in the minimum number of servers. Such migration can be either *hot* or *cold*. Hot (or “live”) migration [6] is performed while the network operation is on going. As such the main concern for hot migration is to minimize the QoS degradation that may be incurred while moving VNFIs. We have proposed a hot migration policy in [7] aiming at minimizing the total cost taking into account both the energy consumption and the reconfiguration costs characterizing the QoS degradation.

This paper deals with migration policies for the case of *cold migration* [8] in which the virtual machines are suspended before performing migrations. This type of migration can be easier to perform, even if it can be quite disruptive for on going traffic. Its application is implemented in cloud environments with redundancy support in which each virtual machine is replicated. The main concern is, therefore, not QoS degradation during migration but rather the *energy consumed during the migration* for transferring the memory associated with migrated virtual machines. The cold migration mechanism is already widely implemented in systems for the management of virtual machines. For instance VMware [9] allows for the cold migration of virtual machines.

This work is an extension of the paper [10], in which we defined and investigated two policies whose aim is the minimization of the total energy. The first one, referred to as Global Policy and based on the application of Markov Decision Process (MDP), allows for the minimization of the total energy in cycle-stationary traffic scenarios and is applicable to the case in which the entire traffic profile is known. The second one, referred to as Local Policy, is applicable in traffic scenarios in which the knowledge of the actual traffic only is assumed. The main differences between the two papers are twofold:

- a new global policy based on the application of the Viterbi algorithm is introduced to minimize the total energy consumption in cycle-stationary traffic scenarios; the advantages of the proposed heuristic based on the Viterbi algorithm allows for a reduction of the computational complexity with respect to the solution based on the Markov Decision Processes and reported in [10];
- an energy-aware VNFI migration problem is formalized and solved for a small network; moreover its performances are compared to the ones obtained with the application of the heuristic based on the Viterbi algorithm;

The paper is organized as follows. Section II introduces the problem of the energy consumption minimization in NFV architectures. The network model and the main assumptions are mentioned in Section III. In Section IV, the Integer Linear Programming (ILP) formulation of the the Energy-Aware

VNF Instance Migration Problem (EVIMP) is reported. The proposed heuristic, based on the Viterbi algorithm is illustrated in Section V. The main numerical results are shown in Section VI. Finally the main conclusions and future research items are discussed in Section VII.

II. PROBLEM OF THE ENERGY CONSUMPTION MINIMIZATION IN NFV ARCHITECTURES

Network Function Virtualization (NFV) involves the instantiation of Virtual Network Function Instances (VNFIs) in order to execute Virtual Network Functions (VNFs) belonging to Service Function Chains (SFCs) [11], [12]. We assume the case in which each VNFI is executing one VNF of a given type (e.g., a virtual firewall, or a load balancer) and the resources (cores, RAM memory,...) assigned to it are shared among more SFCs. The problem of SFC routing and allocation of resources to VNFIs has been investigated in [5] where the formulation of the optimization problem is given. Due to the complexity of the problem, the Maximizing the Accepted SFC Requests Number (MASRN) heuristic has been proposed whose outputs are the routing of SFC and the resource dimensioning of the Virtual Machines associated to the VNFIs. At the end of the application of the MASRN, a logical network is identified in which the nodes are the VNFIs and the links identify which VNFIs are interconnected. This logical network is embedded in the physical network and the MASRN establishes in which physical paths the logical links are routed.

We show in Fig. 1 an example of SFC routing and resource dimensioning. A three levels datacenter network [13] composed by eight servers, eight switches and two routers is considered. Two SFCs are considered in Fig. 1.a. The first one shows that packets need to be processed through two functions: Firewall and VPN Encryption, in that order. The second one involves the execution of a Firewall function only.

To support the two SFCs, two VNFI are instantiated by activating the VMs VM_1^{FW} and VM_8^{EV} in the servers S_1 and S_8 respectively. We illustrate in Fig. 1.b the routing of the two SFCs considered. The first one uses the memory and processing resources of VM_1^{FW} and VM_8^{EV} for the running of the FW and EV VNFs respectively. The second one uses the resources of VM_1^{FW} only. The memory amount and the number of cores allocated to the VMs are also shown in Fig. 1.b. In particular two and three cores have been allocated for VM_1^{FW} and VM_8^{EV} while we assume that the number of cores available in each server is four.

In order to save energy, server consolidation procedures can be activated in low traffic periods when the needed bandwidths and processing resources decreases. As shown in Fig. 1.c, this decrease may lead to the need to allocate only one and two vcores to VM_1^{FW} and VM_8^{EV} respectively. We illustrate in Fig. 1.d how the migration of the VM_1^{FW} in server S_8 leads to the advantage of switching off the server S_1 .

VNFI consolidation as shown in the above example, allows for energy consumption saving, unfortunately it involves

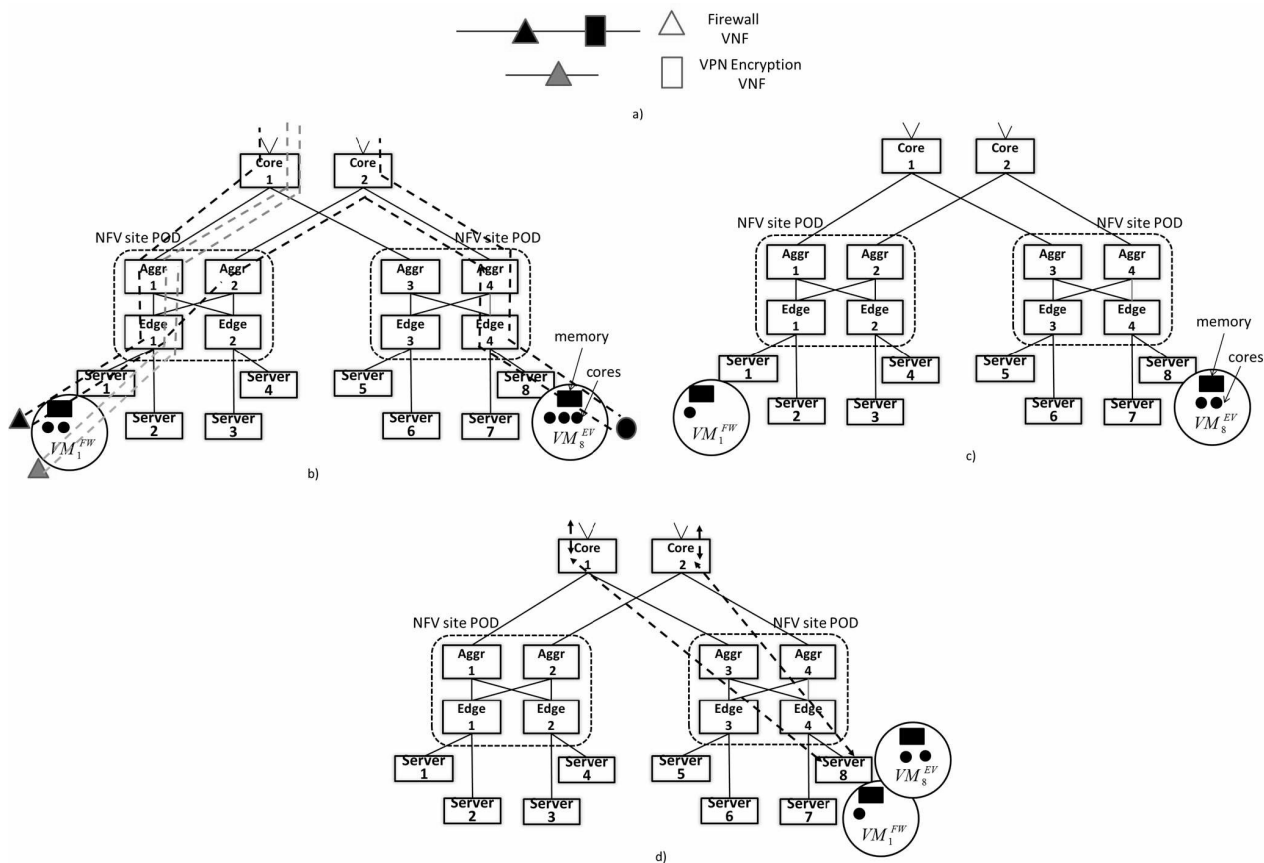


FIGURE 1. An example of SFC routing and resource dimensioning for two SFCs (a); a three levels datacenter network composed by eight servers, eight switches and two routers is considered; memory and cores allocated for the virtual machines VM_1^{FW} and VM_8^{EV} are shown in the peak hour interval (b) and when the traffic decreases (c); the consolidation application leads to the migration of VM_1^{FW} in server S_8 (d).

reconfiguration costs. For this reason sometimes migration could not be recommended or only some migrations should be performed. The investigation of migration policy has been carried out in [7] in the case of hot migration of Virtual Machines. In this case the downtime is minimized with successive stop and copy operations [6] and the main objective is to minimize the QoS degradation due to the bit loss occurring during the downtime.

In this paper we investigate policies in the case of cold migrations where the virtual machine is suspended before performing the migration and high migration times are involved; in this scenario QoS degradation is not the main concern because this type of migration is only applied in case in which the virtual machines are redundant and the switching off of one of them does not impact on the Quality of Service. Conversely the main concern in this scenario is the energy consumption due to transferring the bit of the memory allocated to the migrating virtual machines. This energy consumption characterizes the reconfiguration costs to be taken into account when migration policies, minimizing the total energy cost, have to be identified.

III. NETWORK AND TRAFFIC MODEL

Next we introduce the main terminology used to represent the physical network, VNF and the SFC traffic request [14].

A. PHYSICAL NETWORK

We represent the physical network as an directed graph $G^{PN} = (V^{PN}, E^{PN})$ where V^{PN} and E^{PN} are the sets of physical nodes and links respectively. The set V^{PN} of nodes is given by the union of the three node sets:

- V_A^{PN} : set of access switches of the network where the SFCs are originating and terminating;
- V_R^{PN} : set of switches of the network except the access switches;
- V_S^{PN} : set of servers nodes.

The server nodes and links are characterized by:

- N_w^{core} : processing capacity of the server node $w \in V_S^{PN}$ in terms of number of cores available;
- N_w^{mem} : memory capacity of the server node $w \in V_S^{PN}$ in terms of number of gigabytes available;
- $C_{i,j}$: bandwidth of the physical link $(i, j) \in E^{PN}$ with $i, j \in V^{PN}$;

B. VNF INSTANCE GRAPH

We assume that F types of VNFs can be provisioned as firewall, IDS, proxy, load balancers, ... We denote with $F^{VNF} = \{f_1, f_2, \dots, f_F\}$ the set of VNF types, f_i being the i -th VNF type. The packet processing time of the VNF f_i is denoted with t_i^{proc} ($i = 1, \dots, F$).

The application of an SFC routing algorithm during the Peak Hour Interval leads to identify a logical graph $G^{VNFI} = (V^{VNFI}, E^{VNFI})$ embedded in the physical graph. The nodes of the set V^{VNFI} are the VNFI and the links of the set E^{VNFI} provide information about the VNFI connected. In particular the set V^{VNFI} is given by the union of V_A^{VNFI} and V_F^{VNFI} denoting the set of access and VNFI nodes respectively. In particular the set V_A^{VNFI} of access nodes is fictitious logical nodes to be mapped in the access nodes of the network in which the SFCs are originating and terminating. The introduced graph will be characterized by parameters whose values are dependent on the traffic state T_h ($h = 0, 1, \dots, N - 1$).

C. TRAFFIC MODEL

Though the optimal problem can be defined in general traffic scenarios, we focus its definition in cycle-stationary traffic conditions. In fact, it is well known that in data-center and backbone networks traffic matrices exhibit strong diurnal patterns and are typically cycle-stationary [15], [16], [17]. We denote with N the number of stationary intervals after which the same traffic characteristic occurs again. We denote the duration of a stationary interval with ΔT and we assume that in the h -th ($h = 0, 1, \dots, N - 1$) stationary interval the traffic state T_h occurs. The state T_h characterizes the traffic in the h -th stationary interval between the Virtual Machines implementing the VNFI. Typically, the number of intervals N is equal to 24 and its duration ΔT is 1 hour.

Next we will assume that the traffic state T_0 corresponds to the Peak Hour Interval (PHI) in which the offered traffic is maximum.

IV. ILP FORMULATION OF THE ENERGY-AWARE VNF INSTANCE MIGRATION PROBLEM (EVIMP)

When traffic reduction occurs, the VNF Instance can be dimensioned with a number of cores lower than the one evaluated during the Peak Hour Interval and consolidation techniques can be applied by migrating VNFI so as to occupy fewer servers and to save power consumption. The migration can be performed when the VNFI are supported by Virtual Machines (VM) but at the price of increasing the energy consumption when the memory associated to the VMs are moved. This migration energy may impact the total energy consumption, thus we need VNFI migration policies that allows to minimize it and, consequently, to find a right compromise between energy saving due to consolidation and the energy consumed for the migration.

The main parameters considered in the EVIMP are:

- $P_h^{VNFI}(v)$: processing capacity, expressed in number of cores, required by the VNFI node $v \in V_F^{VNFI}$ in the traffic state T_h ($h = 0, 1, \dots, N - 1$);
- M_v : memory amount, expressed in number of gigabytes, allocated to VNFI node $v \in V_F^{VNFI}$
- $B_h^{VNFI}(d)$: bandwidth requested by the logical link $d \in E^{VNFI}$ in the traffic state T_h ($h = 0, 1, \dots, N - 1$);

- $S_h^{VNFI}(v) = \sum_{d \in I(v)} B_h^{VNFI}(d)$: sum of the incoming link bandwidths in the VNFI node $v \in V_F^{VNFI}$ in the traffic state T_h ($h = 0, 1, \dots, N - 1$); $I(v)$ denotes the set of incoming links to the node $v \in V_F^{VNFI}$;
- $a^{VNFI}(d)$ and $b^{VNFI}(d)$: origin and destination nodes of the virtual link $d \in E^{VNFI}$;
- ϵ_{vw} , it assumes the value 1 if the access node $v \in V_A^{VNFI}$ has to be mapped in the physical access node $w \in V_A^{PN}$; otherwise its value is 0.

The solution to the optimization problem is characterized by two types of variables: the first ones are related to the mapping of the VNFI graph in the network graph of the cycle-stationary period, whereas the second group is formed by the variables related to the server state and the VNFI migration.

The variables of the first group are:

- λ_{vw}^h , binary variable assuming the value 1 if the VNFI node v is embedded in the server node w in the state T_h ($h = 0, 1, \dots, N - 1$); otherwise its value is zero;
- $\tau_{ij}^{h,d}$, binary variable assuming the value 1 if the virtual link d is routed on the physical network link (i, j) in the state T_h ($h = 0, 1, \dots, N - 1$); otherwise it is zero.

The variables of the second group are:

- σ_w^h , binary variable assuming the value 1 if the server node w is turned on in the state T_h ($h = 0, 1, \dots, N - 1$); otherwise its value is zero;
- ξ_v^h , binary variable assuming the value 1 if the node v is migrating when a state change from $T_{(h-1) \bmod N}$ to T_h occurs; otherwise its value is zero;
- ρ_w^h , binary variable assuming the value 1 if the server node w is switching off when the traffic state changes from $T_{(h-1) \bmod N}$ to T_h ; otherwise its value is zero;

The introduced variables have to satisfy the following constraints of the optimization problem:

$$\sum_{w \in V_S^{PN}} \lambda_{vw}^h = 1 \quad h \in [0, \dots, N - 1] \quad v \in V_F^{VNFI} \quad (1)$$

$$\sum_{(i,j) \in E^{PN}} \tau_{ij}^{h,d} = 1 \quad h \in [0, \dots, N - 1] \quad d \in E^{VNFI} \quad (2)$$

$$\begin{aligned} \sum_{i \in \Omega_j^{(-)}} \tau_{ij}^{h,d} - \sum_{y \in \Omega_j^{(+)}} \tau_{jy}^{h,d} &= \begin{cases} \lambda_{b^{VNFI}(d)j}^h - \lambda_{a^{VNFI}(d)j}^h & \text{if } j \in V_S^{PN} \\ \epsilon_{b^{VNFI}(d)j} - \epsilon_{a^{VNFI}(d)j} & \text{if } j \in V_A^{PN} \\ 0 & \text{if } j \in V_R^{PN} \end{cases} \\ & \quad h \in [0, \dots, N - 1] \quad d \in E^{VNFI} \quad j \in V^{PN} \end{aligned} \quad (3)$$

$$\sum_{v \in V_S^{VNFI}} \lambda_{vw}^h P_h^{VNFI}(v) \leq N_w^{core} \quad h \in [0, \dots, N - 1] \quad w \in V_S^{PN} \quad (4)$$

$$\sum_{v \in V_S^{VNFI}} \lambda_{vw}^h M_v \leq N_w^{mem} \quad h \in [0, \dots, N - 1] \quad w \in V_S^{PN} \quad (5)$$

$$\sum_{d \in E^{VNFI}} B_h^{VNFI}(d) \tau_{ij}^{h,d} \leq C_{ij}$$

$$h \in [0, \dots, N-1] \quad (i, j) \in E^{PN} \quad i, j \in V^{PN}$$
(6)

$$\sum_{v \in V_S^{VNFI}} \lambda_{vw}^h \leq U \sigma_w^h$$

$$h \in [0, \dots, N-1]$$

$$w \in V_S^{PN} \quad U \text{ large constant}$$
(7)

$$\xi_v^h \geq \lambda_{vw}^h - \lambda_{vw}^{(h+1) \bmod N}$$

$$h \in [0, \dots, N-1] \quad w \in V_S^{PN} \quad v \in V_F^{VNFI}$$
(8)

$$\rho_w^h \geq \sigma_w^h - \sigma_w^{(h+1) \bmod N}$$

$$h \in [0, \dots, N-1] \quad w \in V_S^{PN} \quad v \in V_F^{VNFI}$$
(9)

In particular constraint (1)-(2) guarantee that each VNFI and each virtual link are assigned to exactly one server node and one network path respectively. Expression (3) establishes conservation flow constraints. Moreover constraints (4)-(5) guarantee that the total number of the cores and the amount of memory occupied by the VNFIs assigned to the server node $w \in V_S^{PN}$ are smaller than or equal to the available number N_w^{core} of cores and the amount N_w^{mem} of memory respectively. Constraint (6) guarantees that the total bandwidth carried by the physical network link $(i, j) \in E^{PN}$ is lower than its capacity C_{ij} . Constraint (7) establishes that any server is turned on when it hosts at least one VNFI. Constraint (8) establishes that a VNFI is migrating when it is not hosted in the same server in two consecutive traffic states $T_{(h-1) \bmod N}$ and T_h . Finally constraint (9) establishes that the server w is switched off when it is turned on in the state $T_{(h-1) \bmod N}$ and turned off in the state T_h .

In the formulation of the Optimal EVIMP, the objective function to be minimized is the total energy consumption E^{tot} in the cycle-stationary interval given by the following expression:

$$E^{tot} = \sum_{h=0}^{N-1} (E_{con}^h + E_{mig}^h)$$
(10)

wherein:

- E_{con}^h is the consolidation energy in the h -th stationary interval; it characterizes the energy consumed by the servers to execute the VNFIs;
- E_{mig}^h is the migration energy during the traffic state change from $T_{(h-1) \bmod N}$ to T_h ; it characterizes the energy consumed to move the memory from a server to another one when any VNFI is migrated.

For the evaluation of E_{con}^h and E_{mig}^h , we assume that all of the servers are equipped with the same amount N_c of cores. Furthermore we also assume a linear model [18] of the server power consumption versus the traffic handled by the server.

A fixed power contribution is also considered when traffic is absent. Hence we assume the following expression of the power consumption P :

$$P = [a + (1 - a) \frac{A_p}{N_c}] P_{max}$$
(11)

wherein P_{max} is the maximum server power consumption, a is the ratio of the baseline power to the maximum power and A_p is the average number of packets handled by the switch.

The energy consumption E_{con}^h is composed by two components: i) the static one $E_{con,f}^h$ represented by the fixed energy consumption related of only server turned on in the states T_h ($h = 0, 1, \dots, N-1$); ii) the dynamic one $E_{con,d}^h$ represented by the variable energy consumption of the VNFIs that depends on the traffic load during the different traffic states.

The assumption of linear server power consumption model allows us to write the following expressions for E_{con}^h

$$E_{con}^h = E_{con,f}^h + E_{con,d}^h$$

$$= \sum_{w \in V_S^{PN}} \lambda_w^h a P_{max} \Delta T$$

$$+ \sum_{w \in V_S^{PN}} \left\{ \frac{\sum_{v \in V_F^{VNFI}} \lambda_{vw}^h S_h^{VNFI}(v) t_v^{proc} / L_p}{N_c} \cdot (\Delta T - D \cdot \xi_v^h) \right\}$$
(12)

where ΔT denotes the duration of a time interval, D is the migration time in which the VNFI is not working, L_p is the packet length and t_v^{proc} is the processing time of the packets handled by the VNFI $v \in V_F^{VNFI}$.

The migration energy E_{mig}^h occurring during the state transition from $T_{(h-1) \bmod N}$ to T_h is due to the moving of memories of VNFIs when these ones are migrated. The migration energy E_{mig}^h could be characterized by the sum of two components: i) the first one $E_{mig,m}^h$ is the energy consumption dependent on the total memory amount that has to be transferred when a mapping change occurs; ii) the second one $E_{mig,f}^h$ is a fix energy component involved when a server is switched off and it is due to the fact that the server has to remain turned on during the migration time D to allow for the migration of the VNFIs migrating. Hence we can write:

$$E_{mig}^h = E_{mig,m}^h + E_{mig,f}^h$$

$$= 2P_{max} \frac{1-a}{N_c} \sum_{v \in V_F^{VNFI}} \xi_v^h \left[\frac{M_v}{L_{agg}} \right] t_{agg}$$

$$+ \sum_{w \in V_S^{PN}} \rho_w^h a P_{max} D$$
(13)

where L_{agg} is the packet length of the flows carrying the memories and t_{agg} is the processing time of the packets containing the memory bits. The term 2 appearing in the expression of $E_{mig,m}^h$ is due to the fact that the migration energy is consumed in both servers involved in the migration.

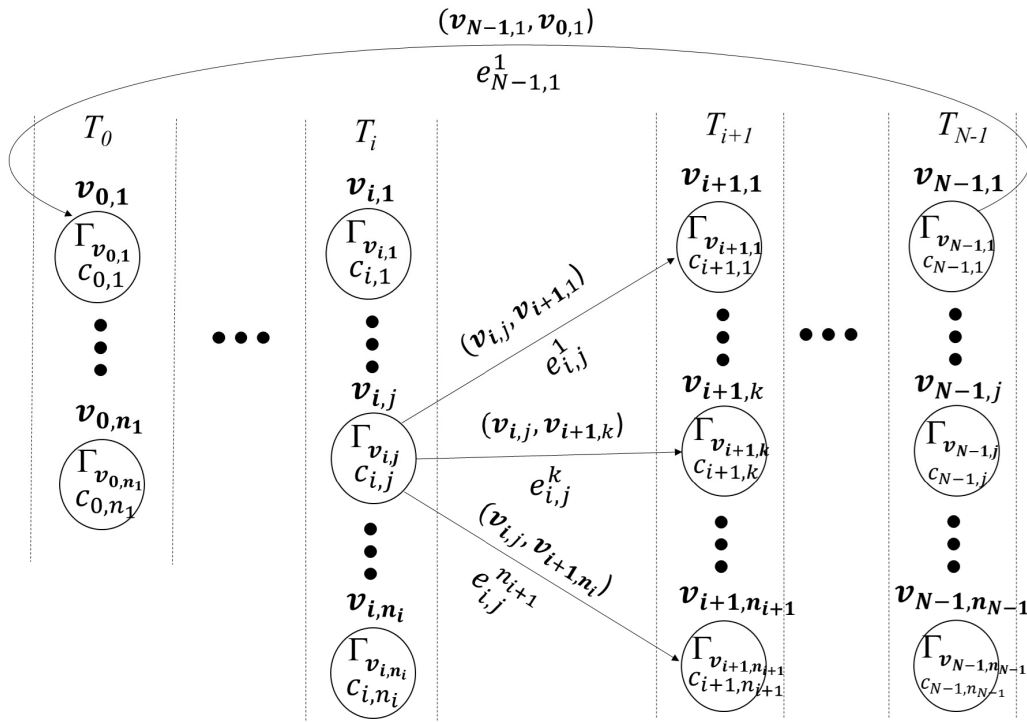


FIGURE 2. The multi-stage graph of the MEAVR heuristic.

To prove that EVIMP is strongly NP-hard it is sufficient to consider the case of one stationary interval ($N = 1$) and infinite link bandwidth. In such a case the EVIMP reduces to the Multi-dimensional Bin Packing Problem that Speitkamp and Bichler [19], Garey and Graham [20] have shown to be strongly NP-hard.

V. ENERGY CONSUMPTION AWARE MIGRATION HEURISTICS

We provide an heuristic whose objective is to determine a *mapping* to be applied in each traffic state. A *mapping* defines in which servers the VNFIs of the set V^{VNFI} are instantiated and in which network paths the virtual links belonging to the set E^{VNFI} are routed. The heuristic, referred to as Migration Energy Aware VNF I Reconfiguration (MEAVR), assumes that the mappings applied in the various traffic states are chosen from an a-priori determined set Ψ . We also assume that the set Ψ is selected containing the mappings $\{\Gamma_i, i = 0, 1, \dots, N - 1\}$, where Γ_i is the mapping minimizing the consolidation energy in the traffic state T_i . We illustrate the MEAVR heuristic in Subsection V-A. Its output is a migration policy \mathcal{D}^{glo} characterized by the set $\Psi^{opt} = \{\Gamma_i^{opt}, i = 0, 1, \dots, N - 1\}$ where Γ_i^{opt} is the mapping to be applied in the traffic state T_i ($i = 0, 1, \dots, N - 1$). The set Ψ^{opt} is chosen so as to minimize the total (consolidation and migration) energy in a cycle-stationary period. We show how the solution of the heuristic consists in evaluating a least cost cyclic path in a multi-stage graph. A low computational complexity methodology to solve the

problem and based on the Viterbi algorithm is described in Subsection V-B.

A. MIGRATION ENERGY AWARE VNF I RECONFIGURATION (MEAVR) HEURISTIC

The determination of the set Ψ^{opt} in the MEAVR heuristic can be accomplished by solving a least cost cyclic path problem in the multi-stage graph \mathcal{G}^{MEAVR} shown in Fig. 2 and organized in N stages numbered from 0 to $N - 1$. The i -th stage contains n_i nodes, whose the generic one is named $v_{i,j}$ ($i = 0, \dots, N - 1; j = 1, \dots, n_i$). The number n_i of nodes in each stage is equal to the number of mappings belonging to the set Ψ and admissible for the traffic state T_i . In particular a mapping is said to be admissible for a traffic state when the instantiation of the VNF I and the routing of the virtual links in the network do not lead to overcome the server processing capacities and the link bandwidth respectively. A node $v_{i,j}$ of the multi-stage graph is characterized by:

- a mapping $\Gamma_{v_{i,j}}$ belonging to the set Ψ and admissible for the traffic state T_i ;
- the consolidation energy $c_{i,j}$ given by the consumed energy by the servers when the mapping $\Gamma_{v_{i,j}}$ is applied;

An edge $(v_{i,j}, v_{(i+1) \bmod N, k})$ of the multi-stage graph characterizes a migration and consequently a mapping change from $\Gamma_{v_{i,j}}$ to $\Gamma_{v_{(i+1) \bmod N, k}}$. The edge is labeled with the migration energy $e_{i,j}^k$.

The values of $c_{i,j}$ and $e_{i,j}^k$ can be determined by the expressions (12) and (13) respectively by appropriately setting the values of the variables $\sigma_w^h, \lambda_{vw}^h$ and ξ_v^h .

The MEAVR heuristic operates according to the following steps:

- the multi-stage graph is built by evaluating the admissible mappings; the consolidation energy $c_{i,j}$ ($i = 0, \dots, N - 1; j = 1, \dots, n_i$) and the migration energy $e_{i,j}^k$ ($i = 0, \dots, N - 1; j = 1, \dots, n_i; k = 1, \dots, n_{(i+1) \bmod N}$) are also evaluated;
- the least cost cyclic path in the graph is evaluated by taking into account that the cost of a path is given by the sum of the consolidation and migration energy of the nodes and edges crossed respectively.

Finally the mappings associated to the nodes of the least cost cyclic path identify the migration policy.

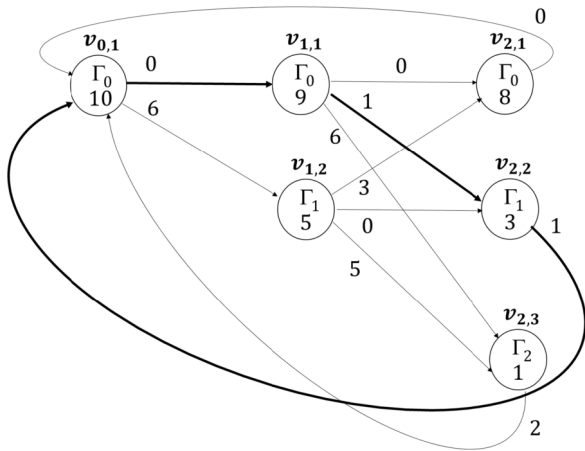


FIGURE 3. An example of application of the MEAVR heuristic.

We show an application of the MEAVR heuristic in Fig. 3. The cycle-stationary period is characterized by three stationary intervals denoted as T_0 , T_1 and T_2 . The mapping minimizing the conservation energy for these states are named Γ_0 , Γ_1 and Γ_2 respectively. We report the multi-stage graph in Fig. 3 in which the number of admissible mappings for the traffic states T_0 , T_1 and T_2 equals 1, 2 and 3 respectively. Inside the nodes $v_{0,0}$, $v_{1,1}$, $v_{1,2}$, $v_{2,1}$, $v_{2,2}$ and $v_{2,3}$ we report the corresponding mapping and the consolidation energy values when the mapping is applied. Each link is labeled with the migration energy needed to perform the corresponding mapping change. Once built the multi-stage graph, the least cost cyclic path can be computed to achieve the migration policy minimizing the total energy in a cycle-stationary period. The path is shown in Fig. 3 with thick lines and it involves a migration policy characterized by the mappings Γ_0 , Γ_0 and Γ_1 to be applied in the states T_0 , T_1 and T_2 respectively. The total energy of the migration policy equals the path cost that is 24.

Finally we conclude by discussing about the complexity of the MEAVR heuristic. Because the determination of a least cost cyclic path can be reduced to the one of a shortest path, we can affirm that the complexity is the one of the Dijkstra algorithm that in the worst case is applied in a graph with N^2 nodes and N^3 links. It follows that if the candidate list in the Dijkstra algorithm is implemented with a binary heap [21],

the computational complexity of the MEAVR heuristic is $O(N^3 \log N)$. This complexity is much lower than the one involved when the migration policy is achieved by applying the Markov Decision Processes [22].

B. SOLUTION OF MEAVR HEURISTIC BASED ON VITERBI ALGORITHM

Usually the peak traffic state (T_0) is characterized by one admissible mapping only (Γ_0); the remaining mappings $\{\Gamma_i, i = 1, \dots, N - 1\}$ of the set Ψ are not admissible in general because they are evaluated in lower traffic conditions and involve the switching off of some servers; this switching off does not allow the support of the peak traffic state without overcoming the server processing capacities and link bandwidth.

When only one mapping is admissible for the traffic state T_0 , a simplified solution can be provided for the MEAVR heuristic. It is based on the application of the Viterbi algorithm [23]. A fictitious stage is introduced in the multi-stage graph of Fig. 4 to build the new multi-stage graph \mathcal{G}_V^{MEAVR} . This stage is composed by one node only denoted as $v_{N,1}$. This node is characterized by the same mapping Γ_0 and by the consolidation energy $c_{0,1}$ of the node $v_{0,1}$. The multi-stage graph is completed by connecting towards the node $v_{N,1}$ those links that in the original graph connect the nodes of the $(N - 1)$ -th stage to the node $v_{0,1}$. According to these changes the multi-stage graph of Fig. 2 becomes the one of Fig. 4. With the introduction of the new multi-stage graph \mathcal{G}_V^{MEAVR} , the problem of determining the policy \mathcal{D}^{glo} consists in finding the shortest path between the nodes $v_{0,1}$ and $v_{N,1}$. Because of the structure of the multi-stage graph \mathcal{G}_V^{MEAVR} , this path can be found by applying the Viterbi algorithm [23] whose the pseudo-code is reported in Algorithm 1. The inputs of the algorithm are: the multi-stage graph

Algorithm 1 Viterbi Algorithm

```

1: Input:  $\mathcal{G}_V^{MEAVR}$ ;  $c_{i,j}$  ( $i = 0, \dots, N; j = 1, \dots, n_i$ );  $e_{i,j}^k$  ( $i = 0, \dots, N - 1; j = 1, \dots, n_i; k = 1, \dots, n_{i+1}$ );
2:  $u_{0,1} = 0$ ;
3: for  $i = 0, 1, \dots, N$  do
4:   for  $j = 1, 2, \dots, n_i$  do
5:      $k_{min} = \arg \min_k (u_{i-1,k} + c_{i-1,k} + e_{i-1,k}^j)$ ;
6:      $\Pi(v_{i,j}) = v_{i-1,k_{min}}$ ;
7:      $u_{i,j} = u_{i-1,k_{min}} + c_{i-1,k_{min}} + e_{i-1,k_{min}}^j$ ;
8:   end for
9: end for
10:  $\mathcal{L} = []$ 
11:  $aux = \Pi(v_{N,1})$ ;
12: for  $i = N - 1, \dots, 0$  do
13:    $\mathcal{L} = [\Gamma_{aux}, \mathcal{L}]$ ;
14:    $aux = \Pi(aux)$ ;
15: end for
16: Output:  $\mathcal{L}$ ;  $u_{N,1}$ 

```

\mathcal{G}_V^{MEAVR} , the consolidation energies $c_{i,j}$ ($i = 0, \dots, N; j = 1, \dots, n_i$) and the migration energies $e_{i,j}^k$ ($i = 0, \dots, N - 1; j = 1, \dots, n_i; k = 1, \dots, n_{i+1}$) (line 1).

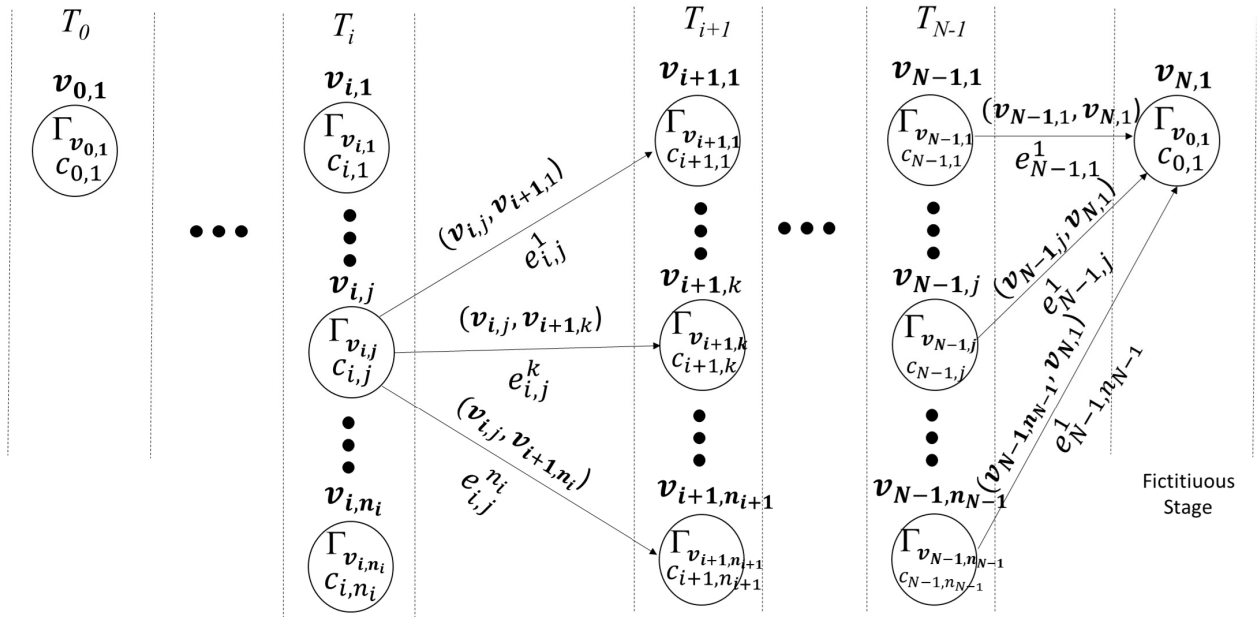


FIGURE 4. The multi-stage graph of the MEAVR heuristic with the addition of a fictitious stage for the application of the Viterbi algorithm.

The cumulative costs $u_{i,j}$ ($i = 0, \dots, N; j = 1, \dots, n_i$) of the least cost path from the node $v_{0,1}$ to the nodes $v_{i,j}$ ($i = 0, \dots, N; j = 1, \dots, n_i$) are evaluated (lines 2-9). Notice that the cost achieved does not include the consolidation cost of the node $v_{i,j}$. In these step of the algorithm, the node before $v_{i,j}$ in the least cost path from $v_{0,1}$ is also evaluated. Such a node is denoted as $\Pi(v_{i,j})$. This operation is needed at the end of the algorithm to identify the nodes of the least cost path from $v_{0,1}$ to $v_{N,1}$. In particular the mappings contained in the nodes of this path characterize the migration policy minimizing the total energy. These mappings are stored in the list \mathcal{L} (lines 10-15).

Finally the outputs of the algorithm are: the list \mathcal{L} of the mappings and the minimized total energy $u_{N,1}$ of the migration policy.

We show in Fig. 3 the application of the Viterbi algorithm for the case of Fig. 5. The dotted arrows are starting from the nodes $v_{i,j}$ ($i = 0, \dots, N; j = 1, \dots, n_i$) and they are directed towards the node $\Pi(v_{i,j})$. The arrows are labeled with the cost of the least cost path from $v_{0,1}$ to $v_{i,j}$. The least cost path from $v_{0,1}$ to $v_{N,1}$ is achieved by following back the dotted arrows starting from the node $v_{N,1}$. In such a way the migration policy minimizing the total energy is characterized by the mappings Γ_0, Γ_0 and Γ_1 with a total energy equal to 24.

Finally we conclude with the complexity of the Viterbi algorithm that is equal to $O(N^3)$ [23].

VI. NUMERICAL RESULTS

We will verify the effectiveness of the VNFI migration policy introduced in Section V. A given number N_{SFC} of SFC requests is generated, each one randomly selected from the

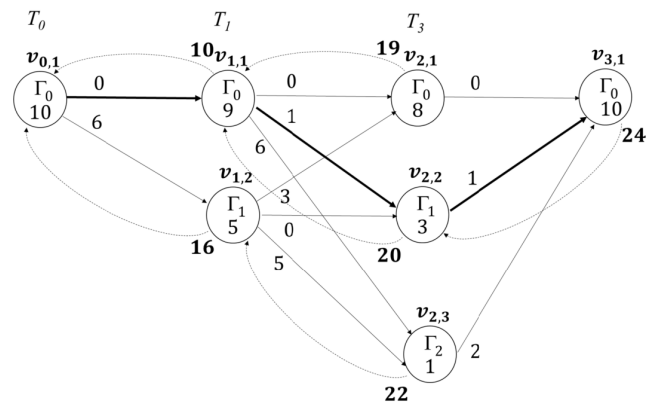


FIGURE 5. An example of application of the Viterbi algorithm.

ones illustrated in Fig. 6: all of SFCs are composed by two access nodes and in addition we can have one Firewall (FW) only (Fig. 6.a), one FW and one Intrusion Detection System (IDS) (Fig. 6.b) and one FW, one IDS and one Encryption VPN (EV) (Fig. 6.c). We also assume that the SFC handles traffic flows of packet length equal to 1500 bytes. The evaluation is carried out according to the data reported in [24] where commercial appliances are used with the values of $120 \mu s$, $160 \mu s$ and $82.76 \mu s$ for the packet processing times of the FW, IDS, EV VNFs respectively. The bandwidth of each SFC in the peak traffic state (T_0) is selected among the values of the set [100Mbps, 150Mbps, 200Mbps, 250Mbps, 300Mbps] according to a Zipf distribution [25]. The SFC routing and the dimensioning of the VNFI in terms of cores are performed during the PHI by using the Maximizing the Accepted SFC Requests Number (MASRN) algorithm proposed in [5].

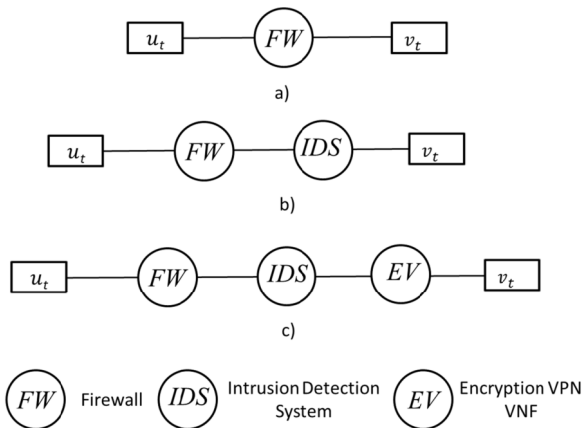


FIGURE 6. Three possible alternatives of SFCs are generated. All of the three SFCs are composed by two access nodes u_t v_t ; the first one is composed by one FW only (a), the second one is composed by one FW and one IDS (b) and the third one is composed by one FW, one IDS and one EV (c).

We evaluate the effectiveness of the introduced VNF migration policies in the case in which a cycle-stationary traffic scenario with N intervals and where the peak bandwidth of the SFCs offered is modulated by the scale factor τ_i in the i -th interval ($i = 0, \dots, N - 1$) chosen according to the classical sinusoidal trend and given by the following expression:

$$\tau_i = \begin{cases} 1 & \text{if } i = 0 \\ 1 - 2\frac{i}{N}(1 - \tau_{min}) & i = 1, \dots, \frac{N}{2} \\ 1 - 2\frac{N-i}{N}(1 - \tau_{min}) & i = \frac{N}{2} + 1, \dots, N - 1 \end{cases} \quad (14)$$

where $\tau_0 = 1$ and $\tau_{\frac{N}{2}} = \tau_{min}$ denote the scale factors in the peak and least traffic conditions respectively.

We assume that the servers are characterized by a maximum power P_{max} equal to 1000 W [7] and an idle power P_{idle} equal to a times the maximum power P_{max} where the parameter a characterizes how much the server power is dependent on the handled traffic; its value can vary from 0 to 1 where $a = 1$ corresponds to the case of server with no rate adaptive power consumption, while $a = 0$ corresponds to that of servers in which the idle power consumption is zero and consequently dependent on the handled traffic only.

The set Ψ of mappings chosen a-priori and introduced in Section V is composed by the elements $\{\Gamma_i, i = 0, \dots, N - 1\}$ where the mapping Γ_i ($i = 0, \dots, N - 1$) is evaluated by applying the VNFI Mapping Minimizing the Power Consumption (VMMPC) consolidation algorithm illustrated in [7] during the interval in which the traffic is characterized by the scale factor τ_i . The VMMPC algorithm, starting from the mapping evaluated by the MASRN algorithm, determines VNFI migrations so as to minimize the number of turned on servers and consequently the power consumption.

First of all we compare the performance of the MEAVR heuristic introduced in Section V to the one achieved by solving the optimization problem introduced in Section IV. The results of the ILP formulation are achieved by using the CPLEX solver. The performance comparison has been

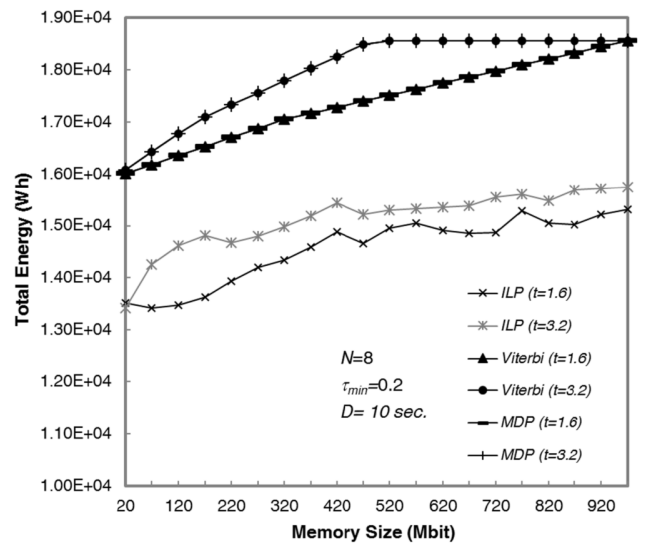


FIGURE 7. Comparison of the results of the MEAVR and Markovian heuristics to the optimal problem. The total energy versus the memory size M is reported for $N_{SFC} = 95$, $N = 8$, $\tau_{min} = 0.2$. The values $1.6 \mu s$ and $3.2 \mu s$ are considered for the memory packet aggregation time.

carried out on a machine characterized by 3.40 GHz Intel i7-3770 processor and by an 8 GB memory. Due to the high computational complexity of the optimal problem, we consider the small network of Fig. 1 composed by eight servers and ten switches. The third level switches are access nodes in which the SFCs are originated and terminated. 40 Gbps links are considered except the links connecting the server whose the bit rate is equal to 10 Gbps. The servers are equipped with 48 cores each one. They are characterized by a value equal to 0.4 of the parameter a . We consider $N_{SFC} = 95$ SFCs each one characterized by the graph reported in Fig. 6.b and composed by two access node, one FW and one IDS. We assume cycle-stationary traffic of parameters $\tau_{min} = 0.2$ and N equal 8. The migration time D is set to 10 sec. We report in Fig. 7 the total energy in a cycle-stationary period as a function of the memory size M to be migrated. The MEAVR heuristic and the optimization results are compared for values of the aggregation time t_{agg} of the memory packet equal to $1.6 \mu s$ and $3.2 \mu s$. Due to the high complexity of the optimization problem, we report a relaxed solution for the CPLEX solver with a percentage error lower than 15%. We also report in 7 the results achieved by applying the heuristic proposed in [7] and based on the application of the Markov Decision Processes. From Fig. 7 we can see how the results of the proposed MEAVR algorithm and the optimization ones are near. For instance in the case of $t_{agg} = 1.6 \mu s$ and $M = 670 Mb$, the total energy values equal 18562 Wh and 17863 Wh respectively. We also verify from Fig. 7 that the results of the MEAVR and Markovian heuristics are the same but with the advantage of a lower computational complexity of the MEAVR heuristic.

Next we evaluate the effectiveness of the migration policy D^{slo} introduced in Section V. We compare it to the following three migration policies:

- local policy \mathcal{D}^{loc} : it is a simple policy based on a cost minimization in each stationary interval and decides the mapping to be applied without taking into account future migration energies; its operation mode is similar to the one proposed in [24]; when traffic state changes from T_i to T_{i+1} the algorithm chooses as mapping to be applied in the state T_{i+1} the admissible one minimizing the difference between the migration energy and the consolidation energy saving involved in a transition from the mapping applied actually in the state T_i ;
- Never Change policy \mathcal{D}^{nc} : it is a policy in which the only mapping Γ_0 is applied and the migration energy is absent; the mapping Γ_0 is the one obtained by the application of the VMPC algorithm during the PHI; it is admissible for all of the traffic conditions and for this reason it is chosen as mapping of the policy \mathcal{D}^{nc} .
- Always Change policy \mathcal{D}^{ac} : it is a policy that applies the least power consumption admissible mapping during each traffic state T_i ($i = 0, 1, \dots, N - 1$); notice how the policy \mathcal{D}^{ac} allows for the minimization of the consolidation energy consumption at the expense of the migration energy.

The performance comparison of the migration policies \mathcal{D}^{glo} , \mathcal{D}^{loc} , \mathcal{D}^{nc} and \mathcal{D}^{ac} are evaluated in a 10-ary fat tree network topology [26]. The network connects 250 servers. 40 Gbps links are considered except the links connecting the server to the rack switches whose the rate is equal to 10 Gbps. The 250 servers are equipped with 48 cores each one. We consider the case in which $N_{SFC} = 2750$ SFCs are generated; the SFCs are originated and terminated from the core switches of the fat tree network.

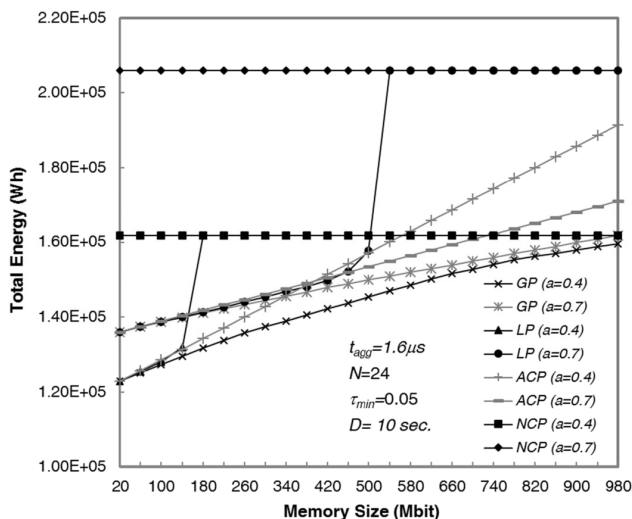


FIGURE 8. The total energy of the policies \mathcal{D}^{ac} , \mathcal{D}^{nc} , \mathcal{D}^{loc} and \mathcal{D}^{glo} as a function of the memory size M associated to each VNFI migrating when $N_{SFC} = 2750$, $N = 24$, $t_{agg} = 1.6 \mu s$ and $D = 10$ sec. Values of the parameters a equal to 0.4 and 0.7 are considered.

We report in Fig. 8 the total energies in a cycle-stationary interval of the policies \mathcal{D}^{glo} , \mathcal{D}^{loc} , \mathcal{D}^{nc} and \mathcal{D}^{ac} as a function of the memory size M associated to each VNFI. We have

assumed the same size M for all of the service functions. The traffic profile is characterized by the factor scales expressed by Eq. (14) in which N and τ_{min} are chosen equal to 24 and 0.05 respectively. The migration time D is set equal to 10 sec, the memory packet length L_{agg} is equal to 1500 bytes and the processing time t_{agg} of each memory packet is assumed to be $1.6 \mu s$. Two sets of curves are reported in Fig. 8 for values of the parameter a equal to 0.4 and 0.7. From Fig. 8 we can observe how the proposed policy \mathcal{D}^{glo} performs better than the policies \mathcal{D}^{loc} , \mathcal{D}^{nc} and \mathcal{D}^{ac} for all of the memory sizes and allows for a minimization of the total energy. For instance for $M = 460 Mb$ and in the case of $a = 0.7$ the total energy values 148953 Wh, 152263 Wh, 152010 Wh, 205889 Wh are obtained for the policies \mathcal{D}^{glo} , \mathcal{D}^{loc} , \mathcal{D}^{ac} and \mathcal{D}^{nc} respectively. As expected the policy \mathcal{D}^{ac} reaches the total energy value of \mathcal{D}^{glo} only for low memory size M . In fact in this case the migration energy is negligible and the optimal policy is the one minimizing the consolidation energy that is \mathcal{D}^{ac} .

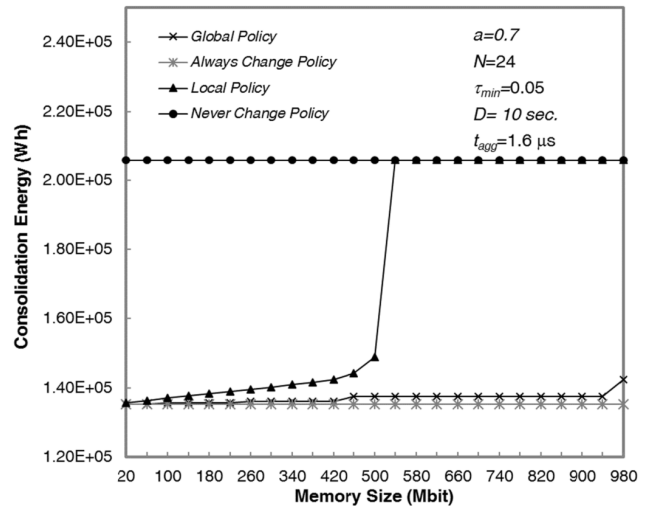


FIGURE 9. The consolidation energy of the policies \mathcal{D}^{ac} , \mathcal{D}^{nc} , \mathcal{D}^{loc} and \mathcal{D}^{glo} as a function of the memory size M associated to each VNFI migrating when $N_{SFC} = 2750$, $N = 24$, $t_{agg} = 1.6 \mu s$ and $D = 10$ sec. The value of the parameter a is chosen equal to 0.7.

We report in Figs 9-10 the consolidation and migration energies for the the case $a = 0.4$ respectively. Fig. 9 confirms how the the policy \mathcal{D}^{ac} minimizes the consolidation energy. This minimization is paid with an increase in migration energy as shown in Fig. 10.

Finally we report in Fig. 11 the percentage energy saving of the policy \mathcal{D}^{glo} with respect to the policy \mathcal{D}^{nc} . This index characterizes the energy saving that the application of a migration energy aware VNFI migration technique allows us to achieve with respect to the case in which migrations of VNFIs is not implemented. We show in Fig. 11 the percentage energy saving as a function of the memory size M for values of the parameter a varying from 0.3 to 1. From Fig. 11 we can do the following remarks: i) the energy saving decrease versus the memory size is due to the increase in energy needed

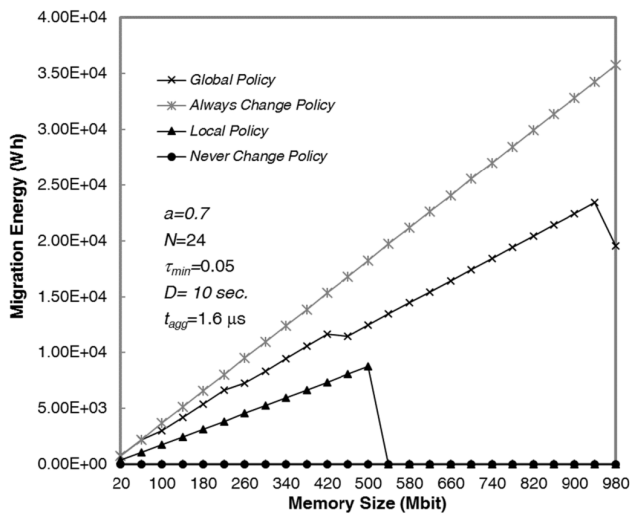


FIGURE 10. The migration energy of the policies D^{ac} , D^{nc} , D^{loc} and D^{glo} as a function of the memory size M associated to each VNF migrating when $N_{SFC} = 2750$, $N = 24$, $t_{agg} = 1.6 \mu s$ and $D = 10$ sec. The value of the parameter a is chosen equal to 0.7.

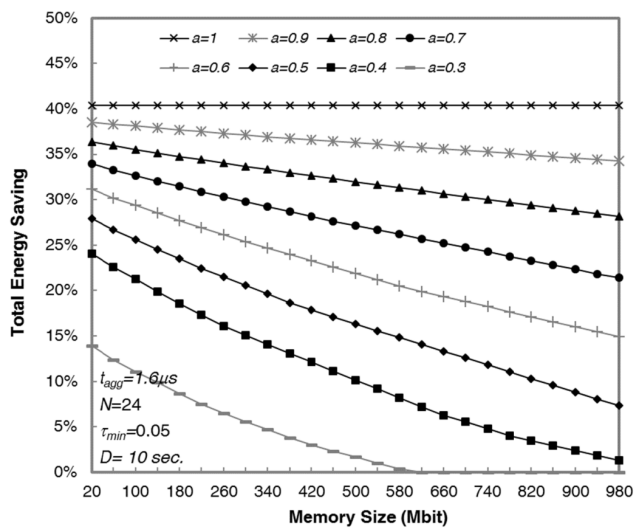


FIGURE 11. The percentage energy saving of the policy D^{glo} with respect to the policy D^{nc} as a function of the memory size M associated to each VNF migrating when $N_{SFC} = 2750$, $N = 24$ and $D = 10$ sec.. The parameter a is varying from 0.3 to 1.

to migrate the VNFs that consequently leads the policy D^{glo} to limit the migrations and to approach the operation mode of the policy D^{nc} ; ii) when the servers are less rate adaptive in power consumption (the value of a is decreasing), the energy saving decreases because the component of fixed power consumption of the servers is smaller and consequently the migration of VNFs is less effective in reducing the power consumption.

VII. CONCLUSIONS

The aim of this paper is to propose migration policies that establish when and where cold migrations of VNFs have to be accomplished so as to minimize the total energy

characterized by the sum of the consolidation and migration energies occurring when the VNFs are moved from the initial location. We have given the ILP formulation of the optimal problem and proposed the simple MEAVR heuristic based on the Viterbi algorithm. We have shown how the MEAVR heuristic is characterized by total energy values near to the optimization results. Finally by applying the MEAVR heuristic we have verified how the proposed migration policy allows for percentage energy saving that can reach the 40% with respect to the solution in which migrations are not performed.

REFERENCES

- [1] L. Gu, S. Tao, D. Zeng, and H. Jin, "Communication cost efficient virtualized network function placement for big data processing," in *Proc. INFOCOM*, Apr. 2016, pp. 604–609.
- [2] R. Mijumbi, J. Serrat, J. Gorricho, N. Bouten, F. De Turck, and R. Boutaba, "Network function virtualization: State-of-the-art and research challenges," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 1, pp. 236–262, 1st Quart., 2016.
- [3] A. Fischer, J. F. Botero, M. T. Beck, H. de Meer, and X. Hesselbach, "Virtual network embedding: A survey," *IEEE Commun. Surveys Tuts.*, vol. 15, no. 4, pp. 1888–1906, 4th Quart., 2014.
- [4] H. Moens and F. De Turck, "VNF-P: A model for efficient placement of virtualized network functions," in *Proc. 10th Int. Conf. Netw. Service Manage. (CNSM)*, Rio de Janeiro, Brazil, Nov. 2014, pp. 418–423.
- [5] V. Eramo, A. Tosti, and E. Miucci, "Server resource dimensioning and routing of service function chain in NFV network architectures," *J. Electr. Comput. Eng.*, vol. 2016, pp. 1–12, Apr. 2016.
- [6] H. Liu, C. Xu, H. Jin, J. Gong, and X. Liao, "Performance and energy modeling for live migration of virtual machines," *Cluster Comput.*, vol. 14, no. 6, pp. 249–264, Jun. 2013.
- [7] V. Eramo, E. Miucci, M. Ammar, and F. Lavacca, "An approach for service function chain routing and virtual function network instance migration in network function virtualization architectures," *IEEE/ACM Trans. Netw.*, to be published.
- [8] V. Medina and J. García, "A survey of migration mechanisms of virtual machines," *J. ACM Comput. Surv.*, vol. 46, no. 3, pp. 1–33, Jan. 2014.
- [9] VMWARE. *Virtualization Overview*, accessed on (2006). [Online]. Available: <https://www.vmware.com/pdf/virtualization.pdf>
- [10] V. Eramo, M. Ammar, and F. Lavacca, "Definition and evaluation of cold migration policies for the minimization of the energy consumption in NFV architectures," in *Proc. ITC*, Genova, Italy, Sep. 2017.
- [11] ETSI. *Industry Specification Group (ISG) NFV*. (Jan. 2015). *ETSI Group Specifications on Network Function Virtualization* [Online]. Available: <http://docbox.etsi.org/ISG/NFV/Open/Published/>
- [12] J. Halpern, and C. Pignataro, Working Group (WG) Service Function Chaining (SFC) Architecture, RFC 7665. (2015). *Internet Engineering Task Force (IETF) Service Function Chaining (SFC)*. [Online]. Available: <https://datatracker.ietf.org/doc/rfc7665/>
- [13] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in *Proc. ACM SIGCOMM Conf. Data Commun.*, New York, NY, USA, 2008, pp. 63–74.
- [14] P. Quinn and J. Guichard, "Service function chaining: Creating a service plane via network service headers," *Computer*, vol. 47, no. 11, pp. 38–44, Nov. 2014.
- [15] S. Gebert, R. Pries, D. Schlosser, and K. Heck, "Internet access traffic measurement and analysis," in *Proc. 4th Int. Conf. Traffic Monitor. Anal.*, Aug. 2012, pp. 40–48.
- [16] G. Maier, A. Feldmann, V. Paxson, and M. Allman, "On dominant characteristics of residential broadband Internet traffic," in *Proc. 9th ACM SIGCOMM Conf. Internet Meas.*, Chicago, IL, USA, Oct. 2009, pp. 90–102.
- [17] M. Z. Shafiq, L. Ji, A. Liu, and J. Wang, "Characterizing and modeling Internet traffic dynamics of cellular devices," in *Proc. ACM SIGMETRICS Joint Int. Conf. Meas. Modeling Comput. Syst.*, San Jose, CA, USA, Jun. 2011, pp. 305–316.

- [18] J. A. Aroca, A. Chatzipapas, A. F. Anta, and V. Mancosor, "A measurement-based characterization of the energy consumption in data center servers," *IEEE J. Sel. Areas Commun.*, vol. 33, no. 12, pp. 2863–2877, Dec. 2015.
- [19] B. Speitkamp and M. Bichler, "A mathematical programming approach for server consolidation problems in virtualized data centers," *IEEE Trans. Serv. Comput.*, vol. 3, no. 4, pp. 266–278, Apr. 2016.
- [20] M. R. Garey, R. Graham, and D. S. Johnson, "Resource constrained scheduling as generalized bin packing," *J. Combinat. Theory*, vol. 21, no. 3, pp. 257–298, 1976.
- [21] V. Eramo, M. Listanti, N. Caione, I. Russo, and G. Gasparro, "Optimization in the shortest path first computation for the routing software GNU zebra," *IEICE Trans. Commun.*, vol. E88-B, no. 6, pp. 2644–2649, 2005.
- [22] M. L. Littman, T. L. Dean, and L. P. Kaelbling, "On the complexity of solving Markov decision problems," in *Proc. 11th Conf. Uncertain. Artif. Intell.*, 1995, pp. 394–402.
- [23] M. F. Bari, S. R. Chowdhury, R. Ahmed, R. Boutaba, O. C. Muniz, and B. Duarte, "Orchestrating virtualized network functions," *IEEE Trans. Netw. Service Manage.*, vol. 13, no. 12, pp. 725–739, Dec. 2016.
- [24] M. Ghaznavi, A. Khan, N. Shahriar, K. Alsubhi, R. Ahmed, and R. Boutaba, "Elastic virtual network function placement," in *Proc. 4th IEEE Conf. Cloud Netw. (CloudNet)*, Oct. 2015, pp. 255–260.
- [25] H. Kikuchi and K. Takahashi, "Zipf distribution model for quantifying risk of RE-identification from trajectory data," in *Proc. 13th Annu. Conf. Privacy, Secur. Trust (PST)*, Jul. 2015, pp. 14–21.
- [26] P.-C. Lin, Y.-D. Lin, C.-Y. Wu, Y.-C. Lai, and Y.-C. Kao, "Balanced service chaining in software-defined networks with network function virtualization," *Computer*, vol. 49, no. 11, pp. 68–76, Nov. 2016.



VINCENZO ERAMO received the degree in electronics engineering and the Dottorato di Ricerca (Ph.D.) degree in information and communications engineering from the Sapienza University of Rome, Italy, in 1995 and 2001, respectively. In 1996, he was a Researcher with the Scuola Superiore Reiss Romoli. In 1997, he joined the Fondazione Ugo Bordoni as a Researcher with the Telecommunication Network Planning Group. From 2002 to 2005, he was an Assistant Professor

and from 2006 to 2010, he was an Aggregate Professor with the INFOCOM Department, Sapienza University of Rome. He is currently an Associate Professor with the Department of Engineering of Information, Electronics and Telecommunications. His research activities have been carried out in the framework of national and international projects. He was a winner of the Exemplary Editor Award 2016 of the IEEE COMMUNICATIONS LETTERS. He was a Scientific Coordinator of the Sapienza University of Rome of E-PhotoOne+ and BONE, two Networks of Excellence focusing on the study of Optical Networks and financed by European Commissions (FP6 and FP7) from 2006 to 2007 and from 2008 to 2011, respectively. He was an Associate Editor of the IEEE TRANSACTIONS ON COMPUTERS from 2011 to 2015. He has been an Associate Editor of the IEEE COMMUNICATION LETTERS since 2014.



MOSTAFA AMMAR (F'02) received the B.S. and M.S. degrees from the Massachusetts Institute of Technology, USA, in 1978 and 1980, respectively, and the Ph.D. degree from the University of Waterloo, ON, Canada, in 1985. He has been with the Georgia Institute of Technology, USA, since 1985, where he is currently a Regents' Professor with the School of Computer Science. His research interests are in network architectures, protocols and services. He has contributions in the areas of multi-

cast communication and services, multimedia streaming, content distribution networks, network simulation, disruption-tolerant networks, virtual network design, and most recently in mobile cloud computing. He has authored extensively in these areas. He has supervised 33 Ph.D. students, many have gone on to distinguished careers in academia and industry. He was a Fellow of the ACM in 2003. He received several awards, including the IBM Faculty Partnership Award in 1996, the Best Paper Award at the 7th WWW conference in 1998, the Georgia Tech Outstanding Doctoral Thesis Advisor Award in 2006, the Outstanding Service Award from the IEEE Technical Committee on Computer Communications in 2010, and the ACM Mobihoc Best Paper Award in 2012. He has served on the research community in multiple roles. He served as the Editor-in-Chief of the IEEE/ACM TRANSACTIONS ON NETWORKING from 1999 to 2003, and he was the Co-TPC Chair of the IEEE ICNP 1997, the ACM CoNEXT 2006, and the ACM SIGMETRICS 2007 conferences.



FRANCESCO GIACINTO LAVACCA received the Laurea (M.Sc.) (*cum laude*) degree in electronic engineering from the Sapienza University of Rome, Italy, in 2013, where he is currently pursuing the Ph.D. degree in information technology engineering with DIET. He has been a Visiting Researcher with the College of Computing, Georgia Institute of Technology, Atlanta, GA, USA. His current research interests are in the fields of all-optical networks and switching architectures, 5G networks, and network function virtualization.

...