Received 15 June 2018; | Revised 6 September 2018; | Accepted 6 December 2018

DOI: xxx/xxxx

SPECIAL ISSUE ARTICLE

# The Pandora of Network Slicing: A Multi-Criteria Analysis

Raphael Vicente Rosa* | Christian Esteve Rothenberg

[1]School of Electrical and Computer Engineering, University of Campinas, Campinas/SP, Brazil

**Correspondence**
*Av. Albert Einstein - 400, Cidade Universitaria Zeferino Vaz, 13083-852. Email: rvrosa@dca.fee.unicamp.br

**Summary**

Evolving requirements of 5G communication business cases around vertical markets in addition to rich content multimedia cloud applications call for advances in enabling technologies and standardization to realize tailor-made end-to-end network services, an approach commonly referred to as network slicing. Different mappings between infrastructure capabilities and custom service requirements produce an overload of information that network service provider must handle to rationalize possible slicing alternatives for a given business case. This article proposes a methodology exploring algorithms from Multi-Criteria Analysis (MCA) to digest a set of slicing criteria advantage points into a coherent set of slicing preferences. Through extensive experimental analysis, we reasoning about observed patterns in the slicing alternatives and discuss the utility of the proposed MCA process for the definition of classes of slices, slicing abstractions views, and recovery/scalability analysis. Altogether, the MCA methodology proposed in this article sheds light into candidate mechanisms to tackle the pandora behind the yet-to-be-understood issues and unanticipated trade-offs of alternative network slicing approaches. Contributions of this article include providing a prominent outlook towards different dimensioning angles of network slicing, scoping its criteria design, timing, and the quality of the criteria information.

**KEYWORDS:**
network slicing, multi-criteria analysis, 5G, SDN, NFV

## 1 | INTRODUCTION

Along the network softwarization path to 5G, network slicing [1] appears as a cornerstone to realize the functional partitioning of infrastructure resources and service capabilities in isolated management/operational scopes to address specific business needs. As examples of network slices, an automotive service (e.g., real-time traffic information) might demand special mobility management and stable broadband support, while a massive IoT service (e.g., electricity metering) might require high scalability disregarding mobility [2]. Traction factors supporting the slice concept include the massive distillery of communication requirements from vertical markets mapped to shared network infrastructure following well-defined abstractions [3]. In simple terms, 5G network slicing contributes to the approximation of applications and networking.

In light of the envisioned 5G services [4], further densification of core and access networks is expected to handle the increase of traffic coming from mobile rich content delivered by cloud applications. Besides, the current infrastructure over-provisioning model offers questionable sustainability to attend upcoming communication requirements of vertical markets (e.g., healthcare, automotive, IoT). Multi-tenancy through broad sharing of diverse types of resource (e.g., network services, functions, compute nodes, etc.) becomes as an desirable characteristic of the emerging 5G ecosystem [5]. A model of Network Slicing as a Service

(NSaaS)[6] presents advantages to operators differentiate data pipes, conceptually via service models and orchestration designated for application, network function, and infrastructure levels. Accordingly, sets of functionality requirements from different network slices would be translated into one or more mappings of infrastructure and network function matched Service Level Agreements (SLAs), associated with service chain links addressing concise traffic characteristics of end-to-end programmable paths and corresponding Quality of Service (QoS) settings.

This article is structured as follows. The next section provides further motivation and discusses relevant related work. Section 3 presents relevant background concepts of network slicing and multi-criteria analysis. Section 4 elucidates the slicing context scoped by this work and presents the proposed multi-criteria analysis approach to network slicing. Section 5 is devoted to the extensive experimental work and results analysis, whereas the Critical discussion of the obtained are associated with research perspectives in Section 6. Finally, the concluding remarks are presented in Section 7.

## 2 | MOTIVATION, RELATED WORK, AND CONTRIBUTIONS

The rise of disaggregation, through new software-centric boundaries and enabling technologies, e.g., advances in Software Defined Networking (SDN), Network Function Virtualization (NFV), Multi-Access Edge Computing (MEC), applied end-to-end (e.g., from radio to core clouds), contributes to the exposure of novel control and management capabilities for network slicing[7]. Aiming at fine-grained network service orchestration[8] and management of end-to-end slices, flexible to comply with diverse SLAs, there exists a challenging demand of comprehensive manners to describe service characteristics, Key Performance Indicators (KPIs), and network element capabilities and requirements[9]. Provided with such a paramount skeleton of programmable artifacts, operators would detain the ability to realize fast network innovation cycles, foremost towards automation, for end-to-end slice deployments. Nevertheless, to establish the suitable manners to handle the large amount and variety of characteristics defining the inventory of active and available slicing features, carriers must incorporate comprehensive methods to digest and rationalize such information for the effective and efficient accomplishment of diverse communication business cases.

A multiple criteria decision making (MCDM) process consists in a system that helps with making decisions under multiple, but conflicting criteria. In the literature,[10] MCDM is commonly divided into Multi-Objective Decision Making (MODM) and Multi-Attribute Decision Making (MADM). For instance, MODM approaches problems such as mathematical programming with multiple objective functions, e.g., optimal embedding of resources over constraining objective functions in NFV[11]. On the other hand, MADM studies the cases where the set of decision alternatives has been predetermined. Being a field of MADM, MCA has been applied in different networking contexts, e.g., assisting the decision process of 4G vertical handovers[12,13], and digesting the diversity and performance of criteria in the offloading paradigm of mobile cloud computing[14].

As a novel approach, in this article we focus on the proposal of a MCA-based methodology capable of dealing with the information overburden behind network slicing design opportunities. As any other outcome from information production, an overload incurs in situations where the decision making process exceeds the capacity to assimilate and act on the information as well as the ability to evaluate every alternative[15]. In general, multi-criteria analysis[16], extensively studied in the literature, including virtual network embedding (VNE)[17], allows flexibility and transparency for accountable recommendations, an appealing approach so far unexplored in the context of network slicing. Considering a set of criteria and certain preferences, through MCA, a large set of network slicing candidate options can be classified, described and ranked to be comprehensively useful in a decision support process (e.g., choices of slices for a particular business case), a model of preferences (e.g., negotiation options in NSaaS offers), and an investigative method for criteria sensitivity (e.g., analysis of slices performance trade-offs). In synthesis, the main contributions of this work consist of: (*i*) the characterization of the dimensioning facets of infrastructure capabilities and business case requirements for network slicing; (*ii*) a methodology for the construction of criteria vectors encompassing slice blueprint descriptors; and (*iii*) an analytic investigation of an MCA methodology in scope of network slices dimensioning.

## 3 | BACKGROUND

### 3.1 | Network Slicing

Scoped by Next Generation Mobile Networks (NGMN)[18], the main problem statements towards 5G uses cases address a high degree of flexibility and scalability demanding network slices in contrast to current monolithic network architectures. In[19], Open Networking Forum (ONF) discusses the applicability of the SDN architecture for 5G slicing particularly to accommodate new and diversified business demands in a cost-efficient way. 3GPP fully incorporates network slicing into the Architecture for the 5G System[20], while in a work-in-progress[21] comprehends the incorporation of network slices into mobile core and radio networks addressing the orchestration and management of 5G services. At Internet Engineering Task Force (IETF), a mailing list discussion named the Common Operations and Management on network Slices (COMS) sparked a series of documents addressing gap analysis, use cases and requirements associated with the existent aspects of Internet protocols to support network slicing[22].

Following the common terminology ideas exposed in most of the standardization bodies, we regard network slicing as an approach to enable multiple logical networks on top of an isolated, fully or partially shared network infrastructure for possibly varied business case purposes[1]. Still, the definitions of how sliced networks will be is yet to be clearly proposed. Slicing decision factors might include: vertical market aims, user service definitions, QoS requirements of each user service, resource availability, etc., which imply in the separation of granular responsibilities among control and data planes according to requirements and capabilities of respective network functions and their execution environments. Such multi-facet decision challenges impact on the key characteristics of slices, such as isolation, shared non-sliced network parts, efficiency on mixed-requirements plus over-engineering of the infrastructure, and the strategy of how operators will expose abstractions concerning intra-domain methods of network slicing (e.g., business case topologies and resources). In[23], aspects of slicing are proposed as part of challenges in realizing slices capabilities, such as:

- **Isolation.** Different slices demand varied degrees of shared infrastructure behaviors, such as interference factors in performance and scalability.
- **Optimization.** The composition of network slices sets available points for better utilization of underlying network resources accommodating different business cases' requirements.
- **Reliability.** In certain aspects, the overall utilization of certain infrastructure footprint also defines assurance boundaries for slices, which might be analyzed according to needed behaviors (e.g., critical applications).
- **Service Mapping.** Based on business requirements, an use case would demand different allocation of resources and respective capabilities.
- **Scaling/Healing.** Dynamics of life cycle management operations towards service maintenance conceives intrinsic analytic opportunities for slices dimensioning.

## 3.2 | Multi-Criteria Analysis

Orthodox decision theory focus on an optimizing approach assuming that different objectives can be expressed with respect to a common denominator by means of trade-offs, meaning the loss in one objective can be evaluated against the gain in another. However, specially for policy enforcement where conflicting interests have to be considered, decision ends up being a multidimensional process. Specific to MCA, in varied decision making support systems, a main role consists in reducing the space of well-determined preferences and options to feasible outcomes, thereby facilitating the resolution of conflicting interests. In such cases, there is no prime solution optimizing all the criteria, so compromise solutions must be found. In addition, preference and indifference relationships among solutions do not represent one outcome might be better than the others, as pairs of solutions might be incomparable with respect to a dominance relation.

MCA consists in elaborating an appropriate set $A$ of potential actions, building a suitable family $F$ of criteria, and determining, for all or some options of $A$, their performances sometimes completed by additional information (possible values for discriminating thresholds, aspiration and/or rejection levels, weights, etc)[15]. A wide variety of algorithms exists in MCA, such as Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS), Simple Additive Weighting (SAW) and Analytic Hierarchy Process (AHP)[15]. Being largely utilized in the literature, the TOPSIS algorithm[24] establishes a distance metric from each alternative from the Positive Ideal Solution (PIS) and to the Negative Ideal Solution (NIS). PIS defines the alternative that maximizes the benefit criteria and minimizes the cost criteria. And NIS is the opposite way, defines the least preferred alternative (i.e., maximizes the cost criteria and minimizes the the benefits). TOPSIS requires the multi-criteria problem to be well-structured (e.g., weights, costs, alternatives, evaluation criteria). Defined PIS and NIS for each criteria, the preference order then is structured according to the relative closeness to PIS utilizing a scalar criterion, which combines these two distance metrics.

# 4 | THE WHY AND HOW OF A MULTI-CRITERIA ANALYSIS METHODOLOGY FOR NETWORK SLICING

## 4.1 | Network Slicing in Context

Flexibility of slicing determines a key enabler for value creation and operational expenses optimization. In varied circumstances, slices are predicted to span diverse technological environments and even different administrative domains[5]. NGMN states network slicing[4] possibly establishing a provider-hosted (e.g., Service Provider) and provider-hosting (e.g., Infrastructure Provider) relationship among administrative domains, where the former express the service requirements and the latter provides the resources/capabilities to fulfill such request. In this context, service providers define mixed requirements and policies (e.g., application regulations, network neutrality liabilities, data privacy guarantees) to infrastructure providers enable NSaaS[6]. Hence, a provider-hosting would expose abstracted views of its resources and slice its infrastructure accordingly, likewise a provider-hosted would only comply with the offered capabilities by provider-hosting, therefore advertising requirements; and when receiving a generic (white-box like) slice of virtualized infrastructure would solve the network embedding problem herself, as provider-hosted understands better its own service than provider-hosting.

In essence, the NGMN[4] establishes a network slice blueprint to define how a network slice instance will be created, referring to the required physical and logical resources for such task. From here on, we utilize the terms network slice and network slice blueprint interchangeably. Based on 5G requirements and foreseen networking use cases (e.g., augmented reality, Internet of Things (IoT), autonomous vehicles), a slice blueprint design and mapping have no simple association with network embedding. Past network embedding approaches consider a fixed set of nodes and links composing well-defined topologies. On the other hand a slice blueprint design would not necessarily define a fixed/strict set of constraints and objective functions for node and link mappings, as it may present varied methods of instantiation. I.e., reasonably in the NSaaS scope, operators would not want to share their slicing methods and underlying fine-granular infrastructure resources – similarly, intra-domain Border Gateway Protocol (BGP) sits well defined only internally to autonomous systems. With the advent of SDN and NFV, both traffic forwarding rules and network functions, realizing a network slice blueprint, can be dynamically placed and chained via various execution environments, and differently by each administrative domain.

## 4.2 | Why a MCA Methodology for Network Slicing?

By expressing diverse requirements for network infrastructure providers, service providers (e.g., vertical market customers) aim to trust in slices determining specific traffic treatments necessary for particular 5G business cases (e.g., eMBB, mIoT and URLLC)[6]. For instance, in case a service provider detains no network footprint in all markets she intends to serve, she might need to partner with access or transit providers to reach their off-net subscribers, consequently being able to explore new roaming scenarios and vertical markets. From an infrastructure provider perspective, the overall allocation of resources to address particular business cases might span varied segments and capabilities of her network, establishing critical factors to scrutinize service providers needs in order to accommodate multiple slices in the same environment. Such aspects include multiple possible deployment options for a slice with diverse technological factors and effects on critical economies of scale. In such NSaaS scenario, a reasonable strategy consists in analyzing the web of relationships among provider-hosting capabilities and provider-hosted requirements as the scope of the slice blueprint design process (see Fig. 1 ).

Our main reason to apply MCA to network slicing concerns the ability to define methods that solve the problems of selecting, describing, comparing, classifying, and ranking a myriad of attributes of multiple slice candidate options in a flexible and effective manner. Abstracted characteristics of providers infrastructure, jointly with possible network function performance profiles and chaining characteristics, define criteria that compose different candidate design options of a network slice blueprint, which can be handled by MCA techniques for different purposes. At the same time, aiming vertical markets and roaming scenarios, service providers need to carefully analyze exposed network slice blueprints as a service, and investigate their varied set of criteria while accommodating their preferences (e.g., costs, SLA, coverage), which scope their pursued business cases. Further decisions, both for service and infrastructure providers, can be taken on top of the set of coherent preferences elucidated by MCA, such as the optimization of objective functions defined by operational or business support systems (OSS/BSS). In a wider scope, the algorithmic issues involving network slicing discussed in[25] represent a coherent view in accordance with our perspective of the slicing design process. I.e., the decision aspects involved in turning a slice blueprint into an instance might contain objective functions (e.g., similar to virtual network embedding[11]) adhering to the policies and resource requirements of the blueprint, and even the infrastructure provider.

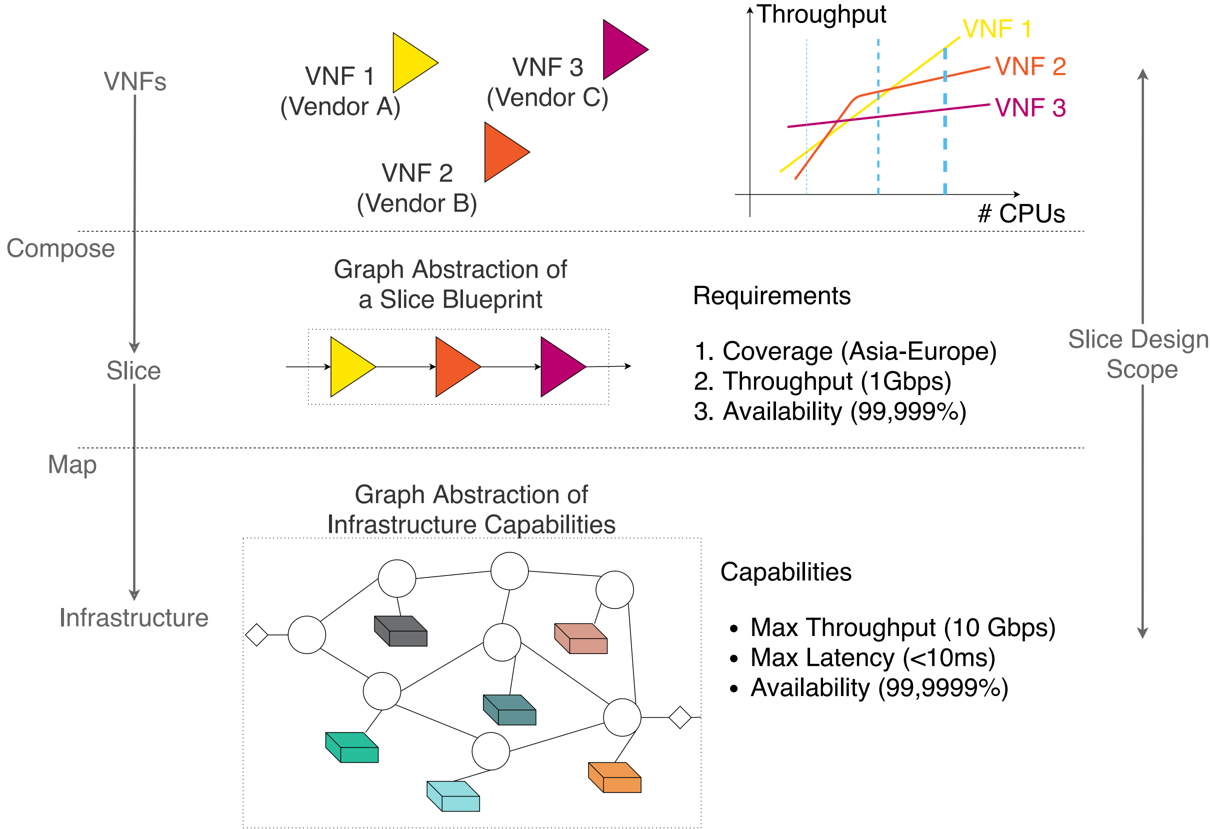


**FIGURE 1** Scope of the slice design process

Envisioned target use cases, potentially benefiting from the applicability of MCA to network slicing, include:

- **Classes of slices.** Incurred segmentations of a infrastructure provider network apt to offer the necessary capabilities to a business case (e.g., chain of VNFs determining SLAs and link characteristics) might define a myriad of possible slice blueprint instances, which can have their criteria carefully examined and classified to represent peculiar characteristics (e.g., performance, fault-tolerance, scalability) and consequent billing options (e.g., Choicenet[26]).
- **Abstraction Views.** The existence of an end-to-end network slice instance can comprehend collaboration among multi-administrative domains at the same time maintaining the fairness for competitiveness, thus requiring minimum information revealed between them while efficient intra-domain traffic engineering to sustain the orchestration of slice blueprints. Inter-domain NSaaS might define a marketplace where administrative domains would expose different abstractions views of their internal infrastructure strategically fine/coarse-grained by varied slicing criteria.
- **Recovery/Scalability Analysis.** While designing a business case mapping, an infrastructure provider might clearly perceive classes of slices by which optional paths might present similar behavior through disjoint infrastructure resources, i.e., opportunities to scale or recover a slice blueprint. When instantiated a previously analyzed slice blueprint, pre-computed differences of configurations and their respective performance impact, would be stored by the orchestration component for life cycle management operations associated with the network slice instance recovery and/or scalability.

## 4.3 | Network Slicing MCA Methodology

In synthesis, a comprehensive analysis of the infrastructure composition of capabilities and the networking business case requirements sets dimensioning angles interfering in possible methodologies that slice blueprints can be designed and mapped onto instances. Based on[16], we elaborate the following steps defining the network slicing MCA methodology explored in this paper:

1. Establish the decision context, containing the aims of the MCA, and the decision makers and other key players – infrastructure and service providers relate as the main actors, possibly in multiple administrative domains, in the scope of NSaaS earlier defined.

2. Identify the options – include the design opportunities a network slice blueprint might detain in order to realize a particular business case.

3. Identify the objectives and criteria that reflect the value associated with the consequences of each option – consists in understanding all the criteria, their technological aspects and importance to the overall process; described in details in 4.3.1.

4. Describe the expected performance of each option against the criteria, possibly scoring the options via the assessment of each their criteria value in face of its consequences – performs the elaboration of possible candidate slice instances enabled by a particular blueprint design; comprehensively detailed in 4.3.2.

5. 'Weighting', assign weights for each of the criteria to reflect their relative importance to the decision – enables the scoring of preferences for each criterion by the involved actor(s) when realizing a particular target use case; discussed in details in 4.3.3.

6. Combine the weights and scores for each of the options to derive and overall value – defines the execution of a chosen MCA algorithm in order to rationalize the slice candidate options and the respective criteria weights.

7. Examine the results – realizes the analysis of the myriad of opportunities and ranking of the slicing candidates.

8. Conduct a sensitivity analysis of the results to changes in scores or weights – executes the MCA algorithm aiming the variability of weights and analyzing the realization of ranked choices coherently with higher/lower criteria weights.

### 4.3.1 | Slicing Criteria

To conceive instantiation options for a particular network slice blueprint and its business case, the representation of a infrastructure must comprehend all the possible capabilities that define critical points in the slice dimensioning angles. Thus, we assume a fully programmable infrastructure consistently exposing its capabilities. I.e., an uniform reference model of a infrastructure provider network represented by a graph as a network topology included compute, storage, and network attributes of links and nodes available for service mappings containing SDN and NFV artifacts. We also assume a network slice blueprint composed as a directed graph by well-defined Virtualized Network Function (VNF) performance profiles distinctively mapping specific allocations of resources (e.g., CPU, Memory, Disk) to packet processing metrics (e.g., latency, throughput, frame loss ratio), jointly with their interconnecting link settings. As object of analysis, we propose a provider infrastructure graph model containing:

- Infrastructure: defined by a node type containing storage and compute resources (e.g., CPU, disk, memory).
- Networks: defined by a node type containing packet processing metrics (e.g., throughput, frame loss ratio, latency) and availability.
- Ports: defined by a node type with specific endpoint property indicating a service termination point in a provider network (e.g., geographic area, edge reachability, another administrative domain endpoint).
- Links: defined by an edge type, annotated its properties (i.e., throughput, frame loss ratio, latency), interconnecting Network nodes to themselves, to Infrastructure nodes, and to Ports.

Aligned with the idea to represent infrastructure provider network capabilities as a graph abstraction, the outcome of a slice blueprint to a slice instance-to-infrastructure mapping contains both technological and topological attributes comprehending the following set of criteria:

- **Performance.** The sum of the impact of the allocation of resources (e.g., CPU, memory, disk) in infrastructure nodes destined to the VNFs of a network slice blueprint. I.e., the slice instance footprint considering its VNFs allocation. For instance, a slice might consume resources from possible overloaded infrastructure nodes, causing a smaller performance impact than in the case of idle ones. This means better performance incurs in less sharing of infrastructure nodes.

- **Betweenness.** The sum of the betweenness centrality of each infrastructure node a slice is mapped to. Betweenness centrality is a graph topological attribute that measures how many all-pairs shortest paths include a node. This gives an indication of the central "monitoring" role played by the node for various pairs of nodes[27].
- **Vitality.** The sum of the closeness vitality of each infrastructure node a slice is mapped to. Closeness vitality is a graph topological attribute defining the change in the sum of distances between all node pairs when excluding that node. This means the importance of a node as a shortest-path connector to all possible paths a graph detains among its nodes.
- **Reachability.** The sum of the closeness centrality of each infrastructure node a slice is mapped to. Closeness centrality is a graph topological attribute that uses the sum of all the distances to rank how central a node is, i.e., it provides a reference to the total distance over all the other nodes, giving the idea of how close a node is to all the other nodes in a graph[27].
- **Cost.** The sum of the placement costs of each VNF along the slice instance mapping.
- **Throughput.** The minimum available end-to-end throughput a slice instance will have when mapped from its input to its output service entering/terminating points.
- **Latency.** The total end-to-end latency a slice instance will have when mapped from its input to its output service entering/terminating points.
- **Hops.** The sum of the hops a slice will have when mapped from its input to its output service entering to terminating points. A hop consists in each node that a path contains in a graph. In the slice case, it will depends on how far/close each VNF is placed to each other and to the input/output slice instance entering/terminating points.

The importance of each listed criteria is defined according to various factors, such as completeness, redundancy, operationality, independence, size, and even timing impacts[16]. Completeness cares if every important criteria was included in the analysis for a proper comparison of the options' performance, meaning no criteria was overlooked and that each one captures key aspects of the objectives concerning the MCA. Redundancy spans the verification of unimportant or duplicated criteria, approaching the assess of meaningful variability for the options to be analyzed, considering a careful analysis of the criteria expressiveness or omission. Operationality concerns how a criteria is defined clearly enough to be assessed, considering a commonly shared and understood scale of measurement. Independence handles the judgment of criteria mutual relationships to recognize that different options satisfy a minimum acceptable level of performance when any given criterion is unaffected in the preference of the others, therefore comprehensively capturing the consequences of the options be independent of each other from one criterion to the next. Size perceives the excessive number of criteria in correspondence with the assess of the MCA results, possibly considering practical importance or inconsistency in the dimension of criteria to prevent the imbalance of the MCA outcome interpretation. And finally, the time factor implies in reviewing the criteria in face of the horizon over which the consequences of the MCA are being valued, i.e., basically refining the impact of the chosen criteria on the options in the short/long-run.

### 4.3.2 | Slicing Options

Given a reference infrastructure network and a slice blueprint referencing a particular business case, the formulation of slicing options consists in structuring possible network slice instances such blueprint might be mapped to while satisfying all its policies and resource requirements. In the methodology proposed by this paper, the definition of how slice instances are designed and have their criteria extracted follows the procedures below:

1. **Abstract the infrastructure.** Consists in the graph modeling of the provider network assets apt to enable NSaaS. In the envisioned context, this step would extract information from all the enabling SDN and NFV technologies realizing programmable service chains able to be mapped to slices. This information would build the infrastructure provider graph model containing infrastructure, network, port and link components. Each one of them annotated with topological attributes and their respective technological capabilities.
2. **Build the Reference Slice Blueprints.** Defining the set of well-formatted business cases supported by the infrastructure, a business case is represented by a slice blueprint reference model, as a direct acyclic graph, that is composed by entering and terminating points adjacent to a sequence of VNFs, each defined by a performance profile requiring specific capabilities.

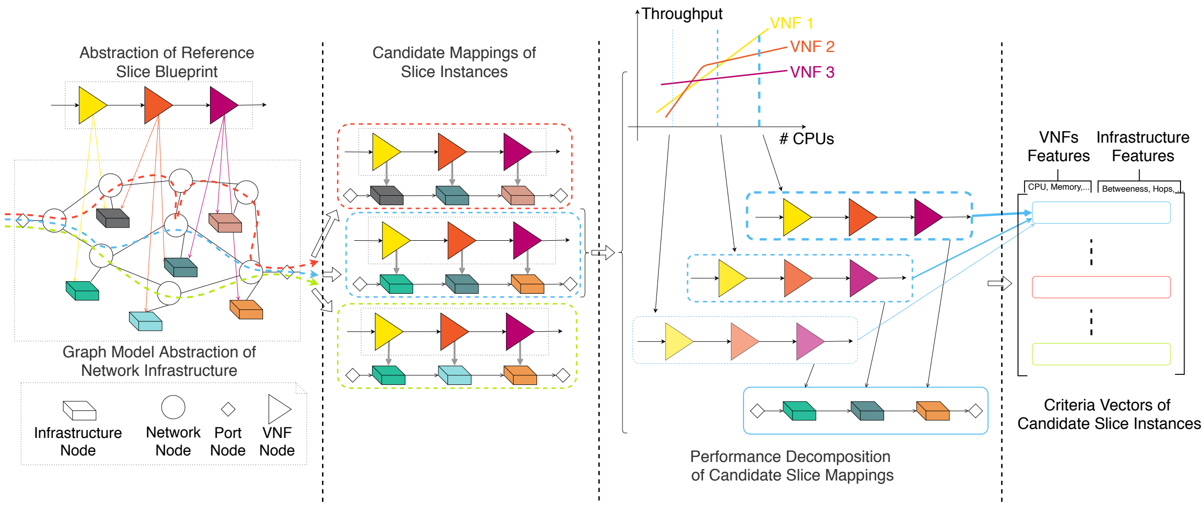


**FIGURE 2** Elaboration of a Candidate Instance Options for a Slice Blueprint

3. **Elaborate Candidate Slice Instances.** Considering the supported business cases, for each one of them is defined an algorithm to elaborate from its reference blueprint the set of possible instance candidates it might be mapped onto. The greedy algorithm, illustrated in Fig. 2 , follows elaborated below:

   (a) Maps the entering and terminating points of the slice blueprint to the required entering and terminating points of the infrastructure graph model.

   (b) For each VNF in the slice blueprint, given its policies and resource requirements, maps it to all possible infrastructure nodes.

   (c) Given all VNF-to-infrastructure mappings, given the slice blueprint link settings and policies, performs the computation of the shortest paths for the slice blueprint adjacencies (i.e., from the entering to the terminating points).

   (d) For each VNF-to-infrastructure mapping in the slice instance, performs the decomposition of all possible optional packet processing metrics a VNF might deliver given its performance profile and the available resources of the mapped infrastructure node. Here, each candidate VNF performance profile decomposition states a possible slice blueprint mapping. Possibly, a single blueprint can be decomposed in a myriad of slice instance mappings.

   (e) Defines all candidate slice instance mappings as vectors of criteria extracted from the end-to-end slice blueprint mapping into infrastructure resources according to the set of criteria previously described.

Through the study of such vectors of criteria it is possible to understand patterns inside a infrastructure according to a particular business case reference blueprint, and elaborate candidate slice instances based on set of its requirements according to the available network infrastructure capabilities. Consequently, given a set of coherent preferences, a MCA algorithm can be applied on the extracted criteria vectors to describe, classify, score, and rank the candidate instance options of a slice blueprint designed for a particular business case.

### 4.3.3 | Slicing Preferences

A critical judgment of the candidate slices could consider the relationship of dominance among them, where one option dominates another if it performs at least as well on all criteria and is strictly better on at least one criterion. Thus, candidate paths would be reduced to a short list, excluded the dominated options. However, these analysis of slicing options might be misleading, as MCA techniques overcome such dominance analysis by a disciplined structure that directs attention to criteria in proportion to the preference they deserve, so called weights. Accordingly, as important as the design of options, the pondered constitution of criteria weights establishes a primary goal in a successful MCA process.

Passed the established sequence of steps to elaborate candidate slices for a single business case, the subject of study now is the arrangement of the candidates according to assigned weights of their criteria. A slice designer role would pursue the

**TABLE 1** Network Infrastructure Profiles

| Profile | Infra Resources | | Link Resources | |
|---|---|---|---|---|
| | CPU | Memory | Throughput | Latency |
| 1 | [2, 4, 8] | [4, 8, 16] | [1, 2, 5, 10] | [6, 8, 10] |
| 2 | [8, 16, 32] | [4, 8, 16, 32] | [5, 10, 20, 40] | [4, 6, 8] |
| 3 | [16, 32, 64] | [16, 32, 64, 128] | [20, 40, 80, 100] | [1, 2, 4] |
| 4 | [2, 4, 8, 16, 32, 64] | [8, 16, 32, 64] | [1, 2, 5, 10, 20, 40, 80, 100] | [1, 2, 4, 6, 8, 10] |

relative importance of each criteria to her analysis, for instance, considering each one of the envisioned target use cases. At different places and in varied situations, an actor in possession of methods to design weight factors for a multi-criteria analysis would endorse preferable conditions that adjust the 'weighting' process to produce coherent slicing preferences. "Once coherent preferences are established, decisions can be taken with more confidence"[16]. Here, the actor stands as a person or a program, both able to update the weights at different timescales according to their target use case and desired set of OSS/BSS preferences (e.g., by automated policy-making based on programmed service intents). These can reflect objectives to coherent decision-making processes aiming, for instance, the power-savings of slices footprint, the priority mapping of reliable low latency slices, or throughput gains for mobile demands of massive broadband services.

## 5 | EXPERIMENTAL EVALUATION

In the absent opportunity to address an actual network infrastructure exposed by a infrastructure provider, in the subsequent evaluated topics we exercise the proposed MCA approach for network slicing via a broad experimental coverage based on the following methodology:

1. **Generate the Random Topology.** A graph defines the set of networks interconnected by links, annotated with throughput and latency attributes. Networks represent one or more programmable set of equipment dedicated to traffic forwarding. Attached to networks, the infrastructure nodes contain CPU and memory capabilities. In the random topology, networks might have ports defining specific service entering and termination points.

2. **Build the Slice Blueprint.** A sequence of ordered VNFs (annotated with CPU and memory requirements) and their interconnecting links (possibly annotated with throughput and latency requirements) define a direct graph concerning a slice blueprint to a specific business case. As input and output port requirements, the blueprint specifically details entering and termination points in the generated random topology.

3. **Elaborate the Candidate Slice Instances.** A simple algorithm realizes the creation of candidate slice instances. For each VNF resources, all the infrastructure nodes are analyzed, the ones with available resources are defined in a set of targets for each VNF. Those sets are shuffled to determine all possible sequences the blueprint can be mapped to the random topology respecting the order of the VNFs. Then, based on the link requirements of the VNFs interconnection, the candidate slices are refined to contain only the suitable infrastructure nodes and links.

4. **Extract the Features and Rank the Slice Instances.** Each of the listed candidate slices, defined by the created blueprint mapping to the random infrastructure topology, have its features processed to construct a vector of criteria, as previously specified (i.e., performance, robustness, vitality, reachability, density, throughput, latency, and hops). In possession of such vectors, the TOPSIS algorithms is applied to the produced slice options, scoring and ranking them.

### 5.1 | Exploring the Experimental Results

A custom software tool-set written in the Python language executes all the previous steps, particularly utilizing some algorithms and graph data structure from the networkx library. Notice, in all the experiments the criteria have abstracted units of measurement. To validate the proposed mechanism, experiments were delineated given the following settings:

**TABLE 2** VNF Profile Metrics

| Cost | CPUs | Memory | Throughput | Latency |
|---|---|---|---|---|
| **0.1** | 2 | 2 | 20 | 20 |
| **0.12** | 4 | 4 | 40 | 16 |
| **0.14** | 6 | 6 | 60 | 14 |
| **0.16** | 8 | 8 | 80 | 12 |
| **0.2** | 10 | 10 | 100 | 10 |
| **0.4** | 12 | 12 | 120 | 8 |
| **0.6** | 14 | 14 | 140 | 6 |
| **0.8** | 16 | 16 | 160 | 4 |

- Infra Model: generated random infrastructures were defined based on the albert-barabasi graph model.
- Infra Model Parameters: each of the infra models had a variable number of nodes defined in 10, 50, and 100, while being created set the neighbour edges equal to 5.
- Compute Nodes Ratio: the amount of infrastructure (compute) nodes added to the random topology was defined in 10%. In further studies we aim to evaluate variations of this rate.
- Compute Nodes Positioning Policy: the compute nodes were linked to network nodes based on vitality, centrality and random policies. For the first two policies, each computed node was added to the topology based on the decreasing value of vitality and centrality of the network nodes it was linked to. And the latter, compute nodes were linked randomly to network nodes. Only one compute node was linked in each network node.
- Infra Profile: four different types of infrastructure profiles were defined according to the set of available resources, shown in Table 1 , respectively addressing infrastructure compute nodes and links with low, medium, high and varied profile of capabilities.
- Slice Blueprint Profile: three different types of reference blueprints were created containing one, two and three VNFs, respectively. The VNF profile is the same for all the reference services, described in Table 2 . For each set of allocated resources a cost value (e.g., monetary, power, operational) was defined according to the VNF performance metrics.

For all the experimental settings the selected range of values expressed in Tables 1 and 2 were defined just for illustrative purposes, and were not based in any related work in the literature, since we could not find any similar reference or similar experiment to the proposed methodology. In each generated infrastructure model two ports were added respectively for the definition of entering and terminating points for the reference blueprint. Such ports were added looking for a pair of network nodes selected with the bigger value of the shortest path distance, calculated based on edges with weights equal to one. Finally, the algorithm to elaborate the candidate slice instances was defined to compute all the combinations of the parameters previously defined. The output dataset of all the slicing candidates for the overall execution created approximately 600 MB of data, turning the experimental validation of the proposed methodology into a data mining job. In the Figures shown and explained below, we dive into the particularities of the dataset.

From Fig. 3 , we have in each sub-figure the distribution of Cost, Latency, and Throughput according to the number of network functions in the slice blueprint, the Infra Profile, and the Infra (Compute Nodes Positioning) Policy. In general, the variations of Cost and Throughput were not affected by variation of the other features. From Fig. 3 (a), we perceive the overall cost of candidate slices increase with the number of network functions in the slice blueprint instances. Similarly, such behavior is seen in Fig. 3 (b), while a pattern among different Infra Profiles show that independently of the number of network functions, the slice blueprint instances latency decreases as the increased amount of available resources in the infrastructure. While in Fig. 3 (c) we do not observe such behaviors, we perceive the overall throughput of the candidate slice instances decreases as the blueprint increases in the number of network functions.

In another point of view, we focused on particular behaviors of the performance metrics of slice instances regarding the following configurations: Infra Profile equals to 3, Infra Policy equals to ‘centrality’, size equals to 100 and network functions (NFs) equals to 3. Both Figs. 4 and 5 show association of Throughput, Latency, Cost, Hops, Betweenness and Vitality metrics.

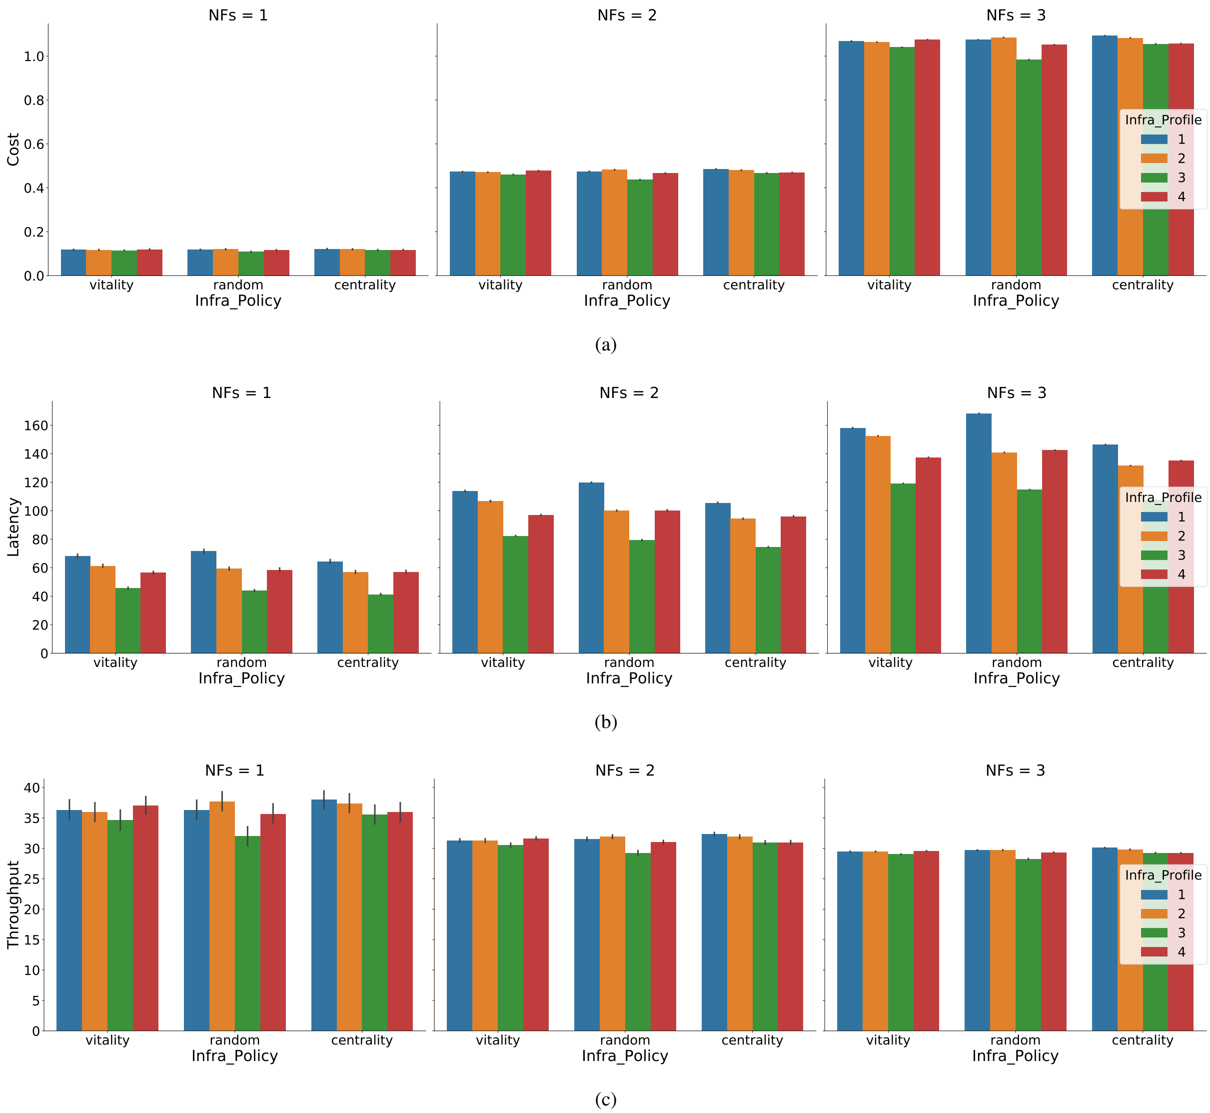


**FIGURE 3** Overall Slicing Candidate Profile Results. (a) Cost (b) Latency (c) Throughput

Mostly important, such figures aim to present the most important patterns of behavior we observed in the dataset for such particular scenario.

According to Fig. 4 (a), there exist clustered slicing candidates detaining the same throughput values and largely varying in Latency, while throughput affects mostly the Cost metric. In another aspect shown in Fig. 4 (a), there exist candidate slice instances clustered in the same amount of Hops, likewise detaining a variety of latency values, while the ones with lower cost have higher values of latency. In what shows Figs. 4 (c) and 4 (d), candidate slice instances with the same amount of hops differ largely in betweenness and vitality metrics, showing not always the ones with lower cost usually were mapped to compute nodes with higher vitality of betweenness. I.e., the results suggest candidate slice instances might vary in trade-off classes of latency and throughput, while detaining long or short paths, and not necessarily being mapped through the most important or central nodes.

The behavioral patterns of classes of candidate slice instances regarding the metrics of latency and throughput become evident when analyzed in contrast with the vitality and betweenness values. As shown in Figs. 5 (a) and 5 (b), clear patterns exist in vertical lines, candidate slice instances with same values of betweenness and vitality, while presenting lower values of cost in association with lower values of latency (i.e., seen by the vertical color patterns, from top-blue to bottom-red). In association with results shown in Figs. 4 (c) and 4 (d), candidate slice instances having the same values of betweenness and vitality are

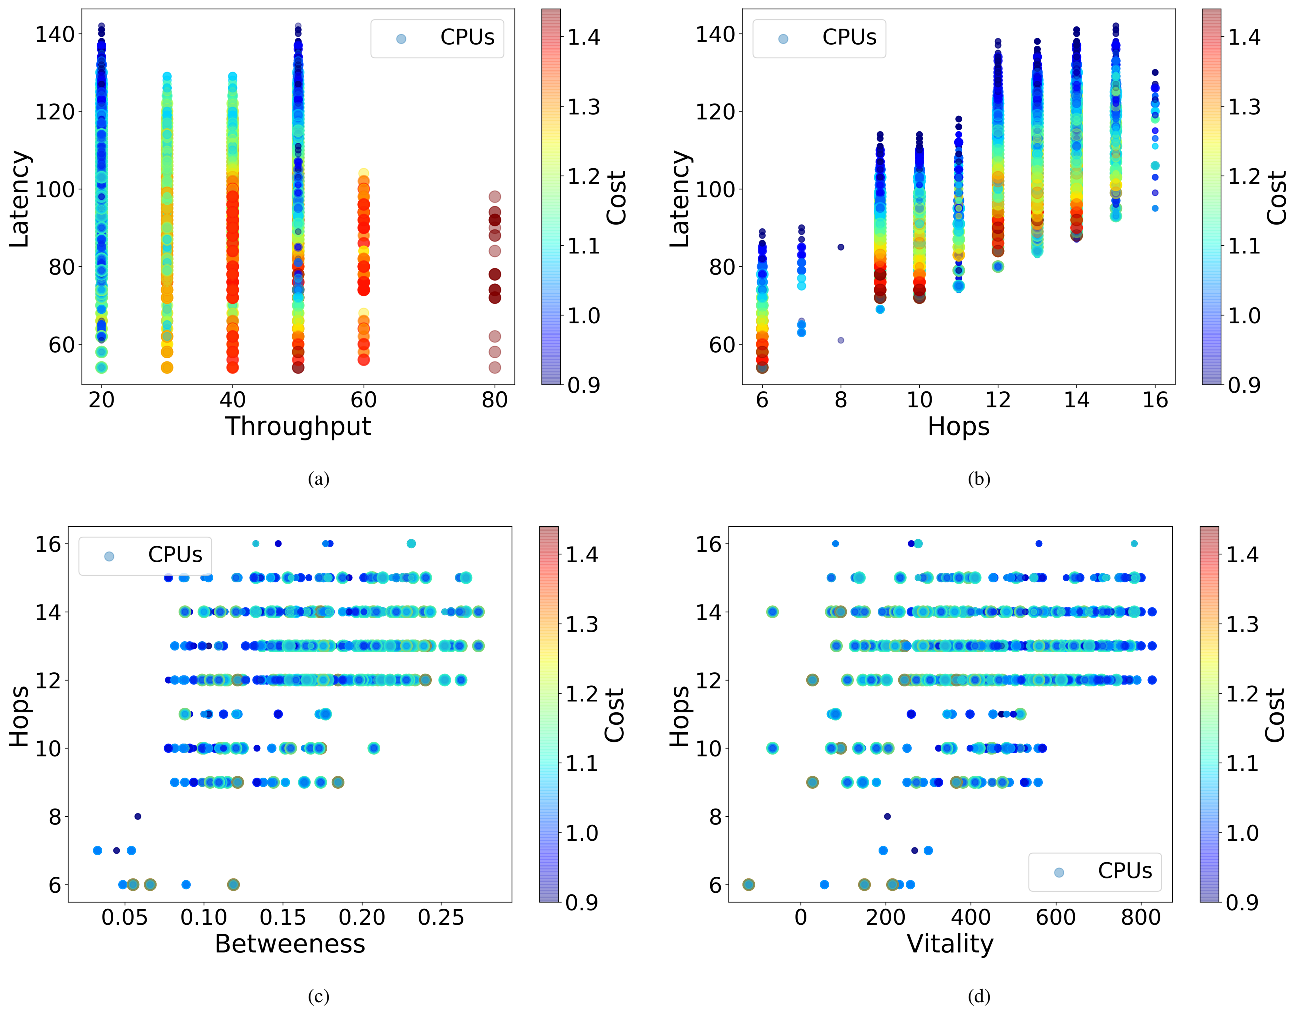


**FIGURE 4** Scatter Plots of Particular Scenario (a) Throughput vs. Latency (b) Latency vs. Hops (c) Betweenness vs. Hops (d) Vitality vs. Hops

characterized by varied sets of hops, and consequently latency. Similarly for the throughput metric, as shown in Figs. 5 (c) and 5 (d), clustered slicing candidates having the same values of vitality and betweenness differ largely in the throughput values, which present the major impact in the cost metric (i.e., the highest cost values are associated with the highest throughput values).

## 5.2 | MCA as a Decision Support System for Network Slicing

Finally, we evaluate and explain how a set of coherent preferences can be used to score candidate slices in order to perceive their criteria and sensitivity to weights of features. Utilizing the algorithm TOPSIS, a set of weights was defined equally for the metrics/criteria Cost, Hops, Latency, Memory, CPUs, Throughput, Betweenness, Reachability, Vitality. Besides, as an input parameter for TOPSIS, some criteria were defined as a benefit (Throughput, Betweenness, Reachability, Vitality) and others as a cost (i.e., Cost, Hops, Latency, Memory, CPUs). For instance, from the execution of the TOPSIS algorithm, a better scored slice candidate would be that one closer to the best values of Throughput while having lower values of Cost and Latency, for instance.

Fig. 6 shows a compilation of examples of trade-off choices when utilizing TOPSIS weights regarding the scenario of features established on the dataset characterized as: Infra Profile equals to 3, Infra (Compute Nodes Placement) Policy equals to centrality, Size equals to 100, and network functions equals to 3. In order to explore the sensitivity of weights in Throughput (benefit) and Latency (cost), three different models of weights were applied upon such particular scenario: defined equally higher priority for Latency and Throughput; higher priority defined for Throughput; and higher priority defined for Latency. A higher

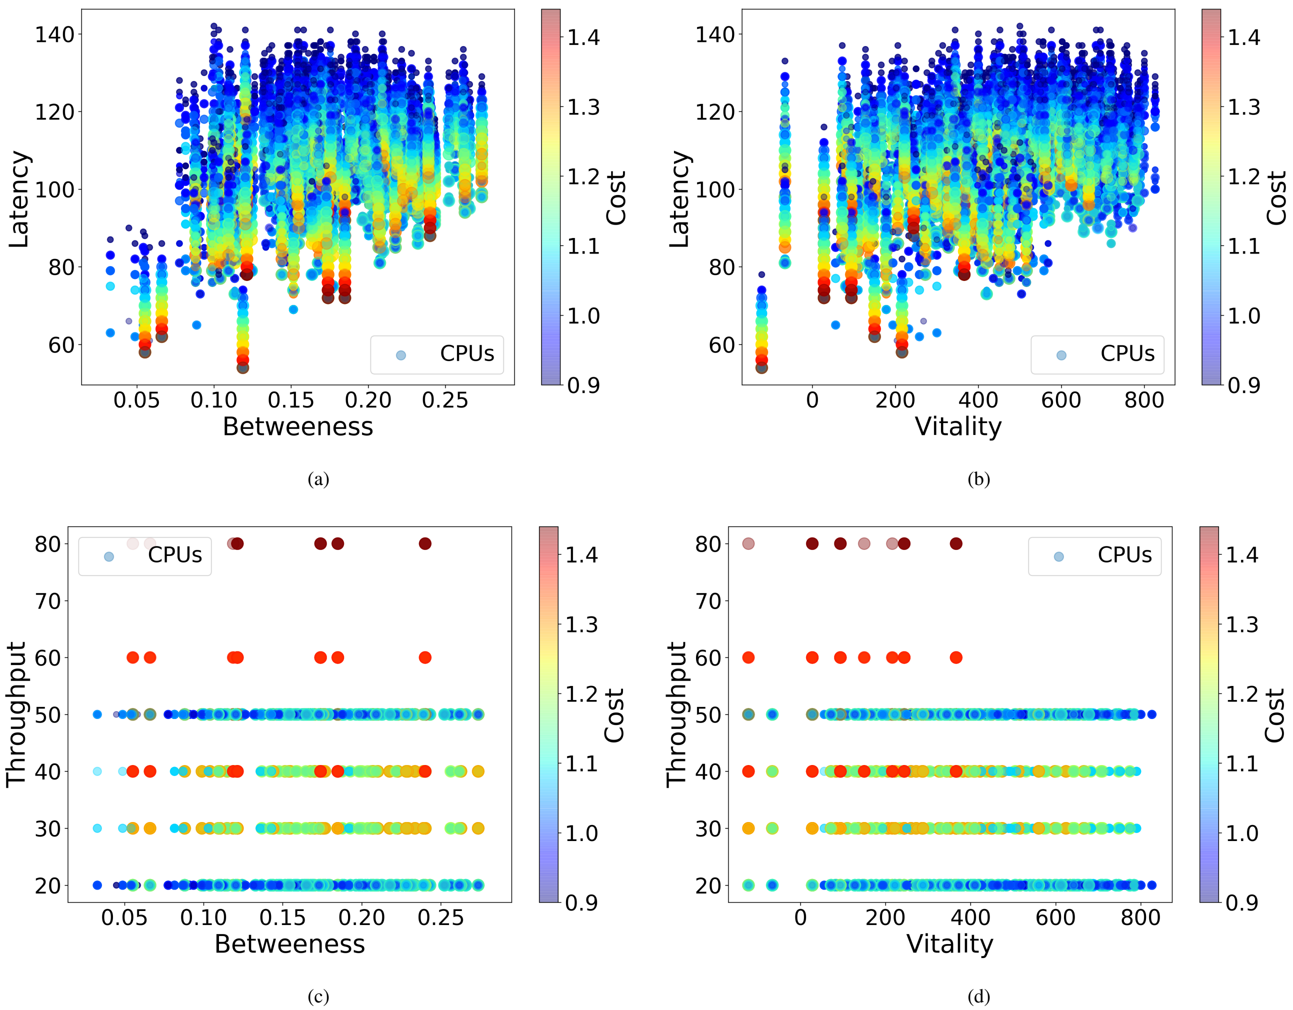


**FIGURE 5** Scatter Plots of Particular Scenario (a) Betweenness vs. Latency (b) Vitality vs. Latency (c) Betweenness vs. Throughput (d) Vitality vs. Throughput

priority means a value of 0.6 in the weight established for a criterion, and all the remaining 0.4 value of weights divided equally between the other criteria.

Evaluating the results in Fig. 6 , we observe the exact trade-offs regarding the choice of the best scores for candidate slice instances when evaluating the Throughput and Latency criteria. In Fig. 6 (a), scored candidate slices present a color pattern, increasing in the score scale, from left-to-right and top-to-bottom as values of Throughput increase and Latency decrease. While such view has an accent distribution of color, while analyzed from the perspective of the Hops criterion, as in Fig. 6 (b), large red circles exist in almost each one of the vertical lines, presenting high scored candidate slices with low latency and high throughput in different path lengths.

While detaining higher priority, the Throughput criterion, as increased in values as a benefit, clearly detain patterns of increased scored values, as shown by Figs. 6 (c) and 6 (d), independent of the latency and or the vitality clusters of candidate slice instances. In opposed aspect, such behavior occurs similarly when higher priority was set to the Latency criterion, while increased in values as a cost presents clear characteristics of lower scored candidate slices. Such evidence is shown in Figs. 6 (e) and 6 (f), when candidate slices detain lowers scores at high Latency values, independently of the Throughput or Vitality criteria.

Regarding the computed scores performed by the algorithm TOPSIS, all the experimental variations of criteria preferences/weights, as shown in Fig. 6 , demonstrate that multiple candidate slices coexist detaining very close scores while having a variety of values in other criteria. Such analysis sustain the clause that depending on the infrastructure provider or customer

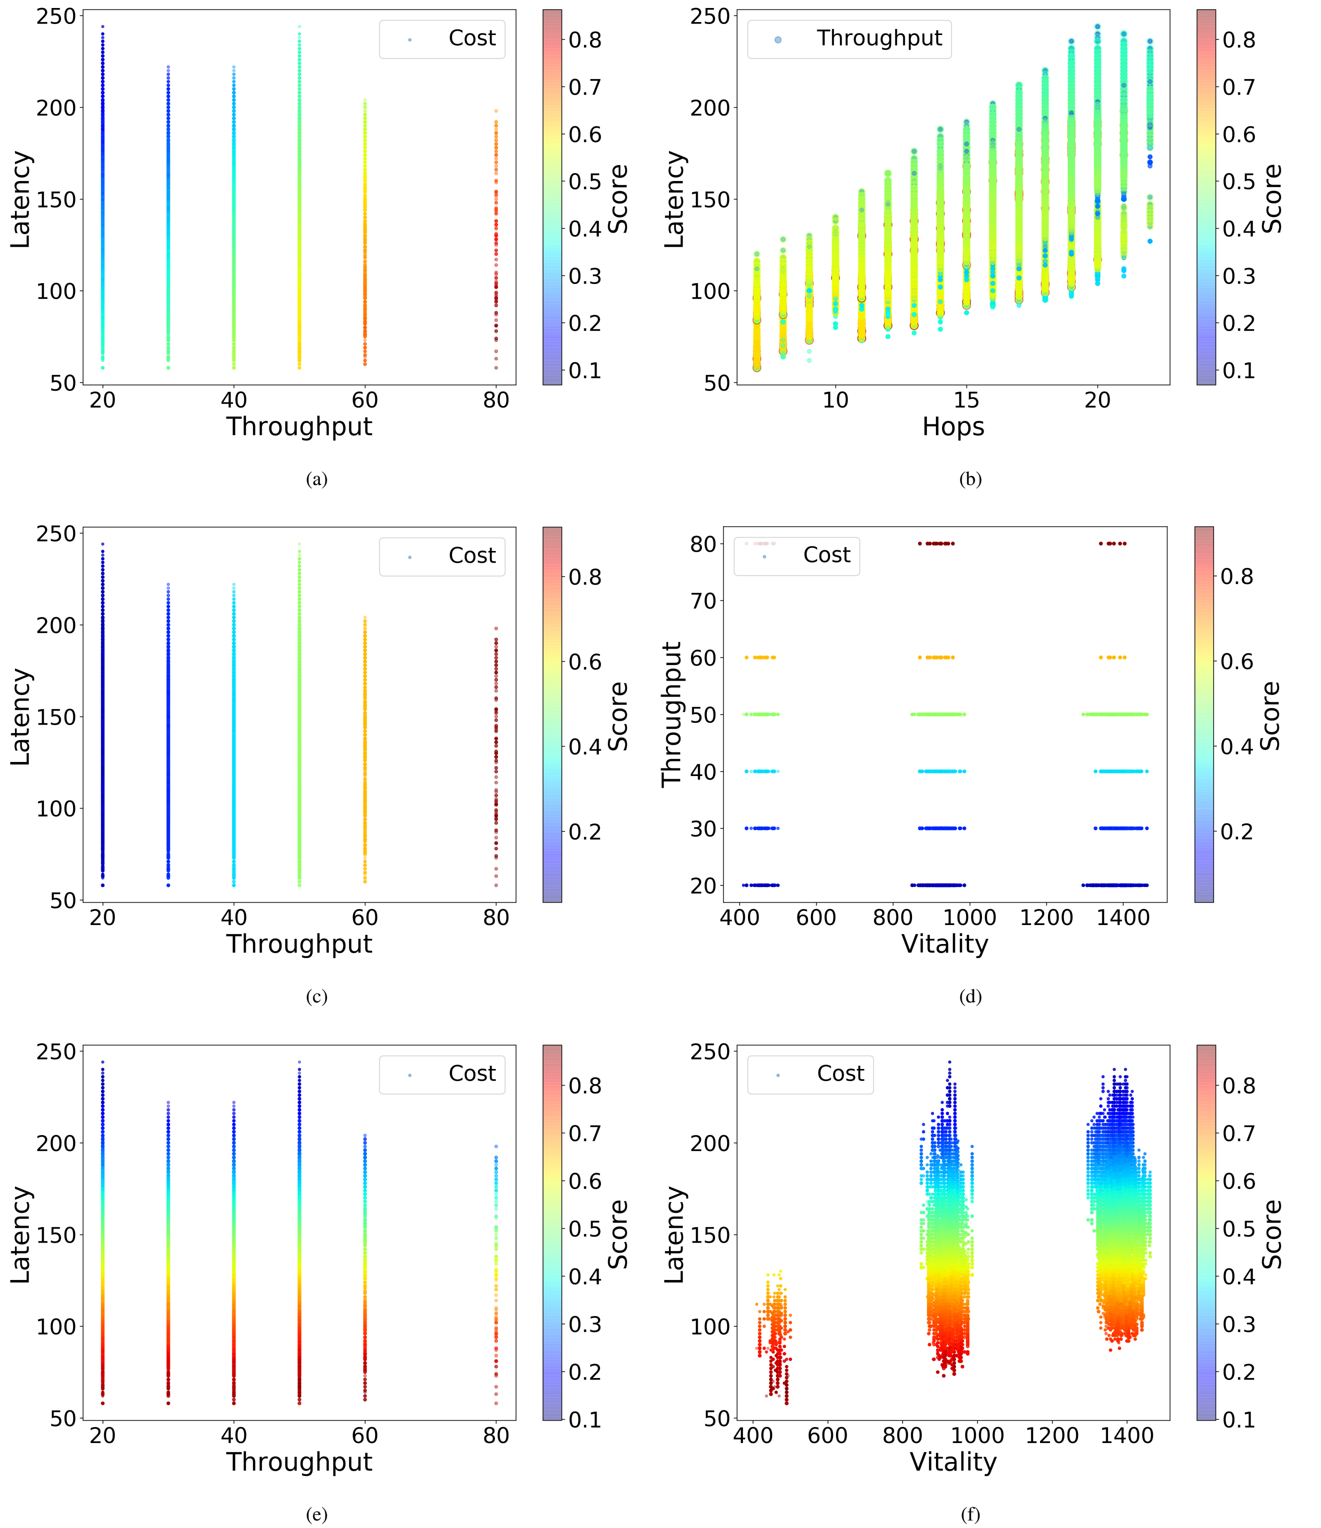


**FIGURE 6** Sensitivity analysis of criteria. (a)-(b) Equal Priority for Latency and Throughput. (c)-(d) Higher Priority for Throughput. (e)-(f) Higher Priority for Latency.

intents for each criteria, classes of candidate slice instances exist in varied settings to be chosen as the means of decision support systems towards the offer, negotiation and realization of a network slicing business case.

# 6 | DISCUSSION

Unlike a top-down optimization problem, by a multi-criteria analysis we aimed to uncover a novel perspective of network slicing. By such means, we quote "thus the concept of 'decision process' has an essential importance. The final outcome is more like a 'creation' than a discovery. With a multiple criteria decision aid the principal aim is not to discover a solution, but to construct or create something which is viewed as useful to an actor taking part in a decision process" (Bernard Roy).

## 6.1 | Results Analysis

Analyzing the experimental results, the utility of a MCA methodology to network slicing can be discussed along different strands.

**Classes of Slices.** Clearly seen in the initial evaluated experiments, a variety of infrastructure capabilities can lead to the differentiation of slice instances, containing specific classes of technological and topological criteria, ranging from high to low scored options. We suggest leading factors in such alternatives relate to the size of the infrastructure and its capabilities as well as the variability of the slice blueprint composition (e.g., number of VNFs and their requirements).

**Slicing Abstraction Views.** Shown by the experiments addressing variable infrastructure profiles, according to the intrinsic aspects of an infrastructure topology and its offered capabilities, the mapping of slice blueprints presents variable instantiation options. The topology characteristics vary based on the way a infrastructure is abstracted, producing different advantage points for each one of the criteria, and even to different business case mapping policies. Thus, we suggest the level of infrastructure abstraction incurs in a view of capabilities that might be refined for different NSaaS purposes.

**Recovery/Scalability Analysis.** Exploring the diversity of slice instance options, via inferred analysis of observed criteria patterns, a slice blueprint mapping can define disjoint paths, which might be utilized as recovery/scalability alternatives to a chosen slice option, and therefore constitute a programmable SLA enhancement. As shown in the experimental results, candidate slice instances might detain similar scores while differentiated in number of hops, likewise similar scores for different amount of hops. This suggests a large set of opportune slice instance mappings with similar characteristics available for recovery or scalability to a single slice blueprint.

**The Complexity of Slicing Options.** The algorithm defined to extract possible slice blueprint mappings detains simplicity, while being greedy and possible costly. The complexity of the algorithm depends on the number of VNFs and their requirements in the slice blueprint, likewise in the size, number of compute nodes, and available resources of the infrastructure graph model. Mostly important, we highlight there might exist a myriad of algorithms (e.g., as stated in[25]) that define candidate slice instances, obeying varied sets of policies, constraints and even objective functions. The design of such algorithms might depend on the target use case. As such, we suggest the involved actors in the MCA methodology to carefully evaluate the variability or specificity of the candidate slice instances generated, and refine the mapping algorithm to comprehend a set of options useful in variety and criteria diversity for a MCA process.

## 6.2 | Inherited Challenges

**Quality of the Slicing Information.** Representing a particular business case, an elaborated slice blueprint references VNFs profiles, which jointly with infrastructure capabilities compose the criteria vectors utilized in the MCA process. In recent work involving the extraction of VNF profiles[28], performance metrics have no strict boundary values, as shown by their confidence intervals. Similarly, the obtained capabilities offered by a infrastructure might fluctuate in time along marginal precision boundaries. We suggest such aspects concern the fuzziness challenges associated with the proposed methodology, which can, and need to, be explored by MCA algorithms that take into account the criteria imprecision in network slice instances.

**Criteria Design.** The extent of the slicing candidate options presented in the experimental analysis depends on the set of criteria established in the design phase involving the reference slice blueprint and the infrastructure graph model. We suggest as long as the elaborated set of criteria respect the well explained considerations in 4.3.1, the reasoning explored in this paper will be maintained when applying the MCA methodology to network slicing. Still, we strongly advocate for a careful examination of the relationships among criteria. Possibly, topological attributes might be directly associated with significant effects in technological aspects of the infrastructure graph model, and consequently affect the quality of the MCA slicing methodology. For those reasons, we suggest an actor involved in the MCA slicing process must be aware of the structural independence, preferential independence, and utility independence of the chosen criteria.

**MCA Issues.** In a first moment, a critical issue about the utilization of MCA algorithms concerns their design choice. Different algorithms share a set of advantages and drawbacks, and a clear methodology to define their choice conceive an important aspect of their performance [29]. For instance, TOPSIS present difficulties to weight criteria and keep consistency of judgment, as it does not consider correlation of criteria. Mainly, MCA algorithms must be analyzed considering their domain of applicability, input of preferences, precision of options values, sensitivity to inconsistent data, consistency of judgment, simplicity, efficiency, and approach to uncertainties. Some approaches even realize a construction of hybrid algorithms, which might also succeed and detain inherent advantages and drawbacks. We suggest a critical way forward consists in defining clear pros and cons for possible algorithms similarly applied to the motivating aspects and in alignment with the methodology proposed by this paper.

### 6.3 | Insights into Future Work

**A View of Promises and Intents.** On the basis of promise theory [30], a slice blueprint states a promise to deliver a set of instantiation criteria under certain performance conditions when conditioned to a specific set of infrastructure capabilities. From this angle, MCA applied to network slicing establishes an opportunity to understand the different promises a slice blueprint mapping might enable towards the criteria approached and their respective weights. On the one hand, the outcome of scored and ranked set of slice candidate instances represent a coherent set of intents a infrastructure provider can assess when addressing a specific slicing business case. On the other hand, the service provider looks for business case instantiation promises by the means of candidate slice instances. Accordingly, we regard intent and promise theories as prominent subjects to make use of the MCA methodology herein proposed.

**Timing Perspectives.** Exploring this paper motivational aspects, the target use cases proposed on the basis of the NSaaS model permit varied slice design ways addressing specific business cases. Given the opportunities of *a priori*, at run-time and *a posteriori* applications of the methodology explained in this paper, some considerations can be addressed. For instance, the joint planning of a slice blueprint and a infrastructure graph model together with classes of slices compose methods to analyze slicing business cases before the deployment stage, and possibly consider conditions regarding what-if policy enforcement to design the utility and cost of candidate slice instances oriented by specific goals (e.g., energy consumption). Moreover, in a continuous analysis of the feedback-loop involving the pre-deployment, run-time and *a posteriori* stages, via the MCA slicing methodology proposed, a recommendation system can be composed to assert and assure only well-scored slicing options.

**Machine Learning Assistance.** In our experimental analysis, as slicing candidate instances were represented by a large set of criteria vectors, we believe such dataset can be rationalized by machine learning algorithms. Through unsupervised learning, we might encounter possible clusters of particular features, correlated criteria, and even outlier slice instances. By supervised learning, we might define classes of a particular criterion labelling candidate slice instances, which might be useful to train regression or classification algorithms. For instance, the regression and prediction results can establish if, given a set of criteria under the performance guarantees of a network infrastructure and a reference slice blueprint, certain slice instances would detain a range of cost or be classified into low latency. Later on, given a subset of network slice instances refined by machine learning algorithms, the MCA methodology proposed in this paper would be realized. However, important aspects must be analyzed. For instance, how marginal errors in regression or the accuracy of predictions affect the MCA process, positively and negatively.

## 7 | CONCLUSIONS

Network softwarization sets a prominent role in network slicing, a way to carriers attain a flexible infrastructure, agility in deploying new customized services, while keeping OPEX/CAPEX under control. Our assumptions concerning network slicing consider well-defined end-to-end slice blueprints containing VNF performance profiles and exposing clear resource requirements that can be translated in a variable set of candidate slice instances depending on infrastructure capabilities. To address the diversity and variety of network slice candidate instances, we propose a MCA methodology as a promising fit to be considered in any networking environment where policy enforcement needs to be applied. Towards 5G, infrastructure providers will need to investigate the return of investments addressing particular business cases, while service provider will need to elaborate adequate strategies to differentiate value and aggregate revenue for upcoming vertical market needs. Based on the experimental results and analysis, we suggest the MCA methodology proposed in this paper can support the network slicing decision processes of infrastructure and service providers. As much as the discussed challenges still impose certain limitations on our approach, we believe the presented insights into future work shed light on new perspectives of applicability of MCA processes into network slicing, specially concerning the perspectives from promise theory and the assistance of machine learning algorithms.

## ACKNOWLEDGEMENTS

This work was partially supported by Ericsson TelecomunicaÃğÃţes (UNI.64) and by the H2020 4th EU-BR Collaborative Call, under the grant agreement no. 777067 (NECOS - Novel Enablers for Cloud Slicing), funded by the European Commission and the Brazilian Ministry of Science, Technology, Innovation, and Communication (MCTIC) through RNP and CTIC.

## References

1. 5GPPP Architecture Working Group . View on 5G Architecture Available at https://5g-ppp.eu/wp-content/uploads/2018/01/5G-PPP-5G-Architecture-White-Paper-Jan-2018-v2.0.pdf – Accessed on 2019-01-05.
2. Afolabi I., Taleb T., Samdanis K., Ksentini A., Flinck H.. Network Slicing and Softwarization: A Survey on Principles, Enabling Technologies, and Solutions. *IEEE Communications Surveys Tutorials.* 2018;20(3):2429-2453.
3. Sharma S., Miller R., Francini A.. A Cloud-Native Approach to 5G Network Slicing. *IEEE Communications Magazine.* 2017;55(8):120-127.
4. NGMN Alliance . 5G Network and Service Management including Orchestration Available at https://www.ngmn.org/uploads/media/170307_5G_Network_and_Service _Management__including_Orchestration_2.12.7.pdf – Accessed on 2019-01-05.
5. Samdanis K., Costa-Perez X., Sciancalepore V.. From network sharing to multi-tenancy: The 5G network slice broker. *IEEE Communications Magazine.* 2016;54(7):32-39.
6. Zhou X., Li R., Chen T., Zhang H.. Network slicing as a service: enabling enterprises' own software-defined cellular networks. *IEEE Communications Magazine.* 2016;54(7):146-153.
7. Ordonez-Lucena J., Ameigeiras P., Lopez D., Ramos-Munoz J. J., Lorca J., Folgueira J.. Network Slicing for 5G with SDN/NFV: Concepts, Architectures, and Challenges. *IEEE Communications Magazine.* 2017;55(5):80-87.
8. Sousa Nathan F. Saraiva, Perez Danny Alex Lachos, Rosa Raphael Vicente, Santos Mateus A. S., Rothenberg Christian Esteve. Network Service Orchestration: A Survey. *CoRR.* 2018;abs/1803.06596.
9. Foukas X., Patounas G., Elmokashfi A., Marina M. K.. Network Slicing in 5G: Survey and Challenges. *IEEE Communications Magazine.* 2017;55(5):94-100.
10. Zimmermann Hans-Jurgen. *Fuzzy Set Theory – and Its Applications.* Springer; 2001.
11. Herrera J. Gil, Botero J. F.. Resource Allocation in NFV: A Comprehensive Survey. *IEEE Transactions on Network and Service Management.* 2016;13(3):518-532.
12. Stevens-Navarro E., Wong V. W. S.. Comparison between Vertical Handoff Decision Algorithms for Heterogeneous Wireless Networks. In: :947-951; 2006.
13. Martinez-Morales J. D., Pineda-Rico U., Stevens-Navarro E.. Performance comparison between MADM algorithms for vertical handoff in 4G networks. 2010;:309-314.
14. Bangui Hind, Ge Mouzhi, Buhnova Barbora, Rakrak Said, Raghay Said, Pitner Tomas. Multi-Criteria Decision Analysis Methods in the Mobile Cloud Offloading Paradigm. *Journal of Sensor and Actuator Networks.* 2017;6(4).
15. Greco Salvatore, Ehrgott Matthias, Figueira Jose Rui. *Multiple Criteria Decision Analysis - State of the Art Surveys.* Springer; 2 ed.2005. International Series in Operations Research & Management Science.

16. JS Dodgson, M Spackman, A Pearman, LD Phillips. *Multi-criteria analysis: a manual*. Available at https://eprints.lse.ac.uk/12761/ – Accessed on 2019-01-05.

17. Zhang Peiying, Yao Haipeng, Qiu Chao, Liu Yunjie. Virtual Network Embedding Using Node Multiple Metrics Based on Simplified ELECTRE Method. *IEEE Access.* 2018;6:37314–37327.

18. NGMN Alliance . Description of Network Slicing Concept Available at https://www.ngmn.org/fileadmin/user_upload/161010_NGMN_Network_Slicing_framework_v1.0.8.pdf – Accessed on 2019-01-05.

19. ONF . SDN Architecture - Issue 1.1 - ONF TR-521 Available at https://3vf60mmveq1g8vzn48q2o71a-wpengine.netdna-ssl.com/wp-content/uploads/2014/10/TR-521_SDN_Architecture_issue_1.1.pdf – Accessed on 2019-01-052016.

20. 3GPP TS 23.501 v2.0.1 . System architecture for the 5G System (5GS) Available at http://www.3gpp.org/DynaReport/23501.htm – Accessed on 2019-01-05.

21. 3GPP TR 28.530 v0.5 - Release 15 . Management of 5G networks and network slicing; Concepts, use cases and requirements Available at http://www.3gpp.org/ftp/Specs/html-info/28530.htm – Accessed on 2019-01-05.

22. Common Operations and Management on network Slices (coms) . Meeting Materials Available at https://datatracker.ietf.org/meeting/101/session/coms – Accessed on 2019-01-05; .

23. Galis Alex. Perspectives on Network Slicing – Towards the New Bread and Butter of Networking and Servicing Available at https://sdn.ieee.org/newsletter/january-2018/perspectives-on-network-slicing-towards-the-new-bread-and-butter-of-networking-and-servicing – Accessed on 2019-01-05.

24. Hwang Ching-Lai, Lai Young-Jou, Liu Ting-Yun. A new approach for multiple objective decision making. *Computers & Operations Research.* 1993;20(8):889 - 899.

25. Vassilaras S., Gkatzikis L., Liakopoulos N., et al. The Algorithmic Aspects of Network Slicing. *IEEE Communications Magazine.* 2017;55(8):112-119.

26. Wolf Tilman, Griffioen James, Calvert Kenneth L., et al. ChoiceNet: Toward an Economy Plane for the Internet. *SIGCOMM Comput. Commun. Rev.*. 2014;44(3):58–65.

27. Zaki Mohammed J., Wagner Meira Jr.. *Data Mining and Analysis - Fundamental Concepts and Algorithms*. Cambridge University Press; 2014. Available at http://www.dataminingbook.info/pmwiki.php/ – Accessed on 2019-01-05.

28. Rosa R. V., Bertoldo C., Rothenberg C. E.. Take Your VNF to the Gym: A Testing Framework for Automated NFV Performance Benchmarking. *IEEE Communications Magazine.* 2017;55(9):110-117.

29. Velasquez Mark, Hester Patrick. An analysis of multi-criteria decision making methods. *International Journal of Operations Research.* 2013;10:56-66.

30. Borril Paul, Burgess Mark, Craw Todd, Dvorkin Mike. A Promise Theory Perspective on Data Networks. *CoRR.* 2014;abs/1405.2627. Available at http://arxiv.org/abs/1405.2627 – Accessed on 2019-01-05.

# AUTHOR BIOGRAPHY

imgs/raphael.png

**Raphael Vicente Rosa.** currently pursues his thesis on multi-domain distributed NFV as a PhD student in University of Campinas, Brazil. During the last two years, he worked as a visiting researcher in Ericsson Research Hungary, where he contributed to EU-FP7 Unify project and developed activities within H2020 5G Exchange project. His main interests sit on state-of-the-art SDN and NFV research topics entangled with disruptive technologies, aiming to open source his pet projects.

imgs/chesteve.png

**Christian Esteve Rothenberg.** is an Assistant Professor in the Faculty of Electrical & Computer Engineering (FEEC) at University of Campinas (UNICAMP), Brazil, where he received his Ph.D. and currently leads the Information & Networking Technologies Research & Innovation Group (INTRIG). His research activities span all layers of distributed systems and network architectures and are often carried in collaboration with industry, resulting in multiple open source projects in SDN and NFV among other scientific results.