

Automated Decision Making for the Multi-objective Optimization Task of Cloud Service Placement

Michael Seufert, Stanislav Lange, Markus Meixner

University of Würzburg, Institute of Computer Science, Chair of Communication Networks, Würzburg, Germany
 {seufert | stanislav.lange | markus.meixner}@informatik.uni-wuerzburg.de

Abstract—The network functions virtualization (NFV) paradigm provides advantages with respect to aspects like flexibility, costs, and scalability of networks. However, management and orchestration of the resulting networks also introduce new challenges. The placement of services and virtualized network functions (VNFs) is a multi-objective optimization task that confronts operators with a multitude of possible solutions that are incomparable among each other. The goal of this work is to investigate mechanisms that enable automated decision making between such multi dimensional solutions. To this end, we investigate techniques from the domain of multi attribute decision making that aggregate the performance of placements to a single numeric score. A comparison between resulting rankings of placements shows that many techniques produce similar results. Hence, placements that achieve good rankings according to many approaches might be viable candidates in the context of automated decision making.

Index Terms—Cloud Service, NFV, Placement, Orchestration, Multi-Objective Optimization.

I. INTRODUCTION

The Network Functions Virtualization (NFV) paradigm offers numerous benefits to network operators in terms of costs, flexibility, scalability, and vendor independence. In contrast to the prevalent deployment of specialized middleboxes for network functions like firewalls or load balancers, NFV leverages virtualization mechanisms in order to perform the packet processing tasks of the former via software that runs on commercial off-the-shelf (COTS) hardware.

However, management and orchestration techniques are required in order to achieve and maintain a high degree of flexibility and assert that QoS and QoE constraints are met. In particular, the placement of virtualized network functions (VNFs) within the network can have a significant impact on both, user and operator satisfaction. Since goals like low latency among VNF instances and low latency between VNFs and end users can be competing, finding suitable VNF placements corresponds to a multi-objective optimization task.

In addition to the increased complexity of algorithms that can solve such problems, the solutions they return can not always be compared with each other due to different domains and units of the objectives. Especially in the context of automated and dynamic service migration and instantiation, however, algorithms need to choose one distinct solution.

The contribution of this work is threefold. First, 4 methods for determining the relative importance of different objectives are presented and compared with each other. In contrast to approaches that determine such weights a priori, the methods

presented in this work take into account characteristics of the solutions that are returned by the multi-objective optimization algorithm. Second, four mechanisms for aggregating the performance of a multi-dimensional solution into a single score are introduced. Finally, the rankings of solutions that result from different combinations of weighting and aggregation techniques are characterized. On the one hand, analyzing solutions that consistently achieve high ranks according to many approaches might lead to more efficient methods for identifying viable placements. On the other hand, the comparison can help derive guidelines for choosing the appropriate ranking mechanism for a particular problem. All comparisons are performed on realistic problem instances featuring graphs from the Internet Topology Zoo [1] and three objectives.

The remainder of this work is structured as follows. After an overview of related work in Section II, the data set is presented alongside the resulting problem instance in Section III. Methods for assessing the weight of each objective dimension are introduced and compared in Section IV. These methods are then used as input for algorithms that assign a score to each placement. In Section V, four such algorithms are discussed and compared with respect to the rankings of placements they produce. Finally, Section VI concludes the work.

II. RELATED WORK

Resource management in clouds [2], in particular, the placement of cloud services in data centers has become an increasingly important problem. Typically, many parameters and metrics regarding resource utilization and performance have to be taken into account within the cloud and the network. Thus, different methodologies are proposed in literature to place the virtual machines efficiently. In the context of placing virtual network functions (VNF), [3] investigates a weighted sum approach, while [4] uses a linear program to find an optimal placement. Also [5], [6] propose linear programs for chains of VNF, while [6] adds a Pareto analysis to investigate the trade-offs between the different dimensions.

A related problem, which has been discussed recently, is the placement of SDN controllers. It is also a multi-objective optimization problem, which has to take into account a large set of parameters and metrics. Weighted sums (e.g., [7]) and linear programs (e.g., [8]) are widely used. Additionally, the Pareto frontier is analyzed when different alternatives are incomparable. Due to state explosion, the problem of obtaining the Pareto frontier is frequently tackled heuristically [9], [10].

However, no automated decisions can be taken from Pareto frontiers, which will be tackled in this work.

Therefore, we will transform the Pareto frontiers into a ranked list of alternatives. To compare the rankings when the underlying order of alternatives is unknown, we will mainly rely on correlation coefficients and techniques based on probabilistic ranking models. In [11], the rank correlation between the pairs of ranking is calculated using either Spearman's ρ or Kendall's τ . [12] proposes a measure of agreement between rankings based on removal of disputable elements. A basic model for order statistics was developed by Thurstone [13], and [14] constructed an equivalent model based on choice probabilities. Mallow [15] presented simplified and analytically tractable models induced by paired comparison. [16] investigates concordance between different judges (i.e., rankings) based on Mallow's model to detect outlier rankings. [17] proposes to compare the distribution of ranks by box plots and derive a degree of discordance based on the interquartile range. The goodness of fit of simple ranking models is investigated in [18], and metric based ranking models are discussed in [19]. A classification of probabilistic ranking models can be found in [20].

III. DATA SET DESCRIPTION

In order to investigate the practical feasibility of the different weighting and ranking methods that are discussed in this work, realistic input data is required. To this end, we use multiple network graphs from the Internet Topology Zoo [1] and evaluate possible service placements with respect to a total of three objective functions. While results are consistent among different networks, some characteristics depend on statistics like the number of nodes and the diameter of the graph. Hence, the Internet2 OS3E topology is chosen as an exemplary representative. Table I provides an overview of the graph as well as the resulting problem instance.

In order to keep the solution space small enough to visually illustrate the effects and behavior of the presented methods, only four service instances are placed in the network. Although this results in a total of 46,376 distinct placements, only ten of those are Pareto optimal and thus relevant during the decision making process. In the context of larger search spaces, e.g., when placing more services or dealing with networks that have more nodes, an exhaustive evaluation of all possible placements might not be feasible due to time and resource constraints. For such cases, a trade-off between accuracy and runtime can be achieved by employing heuristic approaches that can approximate the Pareto frontier [9].

As mentioned in the previous paragraph, three different objective functions are taken into account when assessing the performance of each placement. These include two latency-related measures, namely, the mean and average latency between services and end users. We use the longitude and latitude information that is provided for each node to calculate the Euclidean distance between nodes and approximate the latency of each link. For multi-hop paths, the latency is defined as the sum of latencies of all involved links. Furthermore, the

TABLE I: Information regarding the network topology and the corresponding problem instance used in this work.

Property	Value
Graph name	Internet2 OS3E
Number of nodes	34
Number of placed services	4
Number of distinct placements	46,376
Number of Pareto optimal placements	10
Objective functions	Mean latency to end users $\pi^{\text{avg latency}}$ Maximum latency to end users $\pi^{\text{max latency}}$ Imbalance between service instances $\pi^{\text{imbalance}}$

load imbalance between service instances is defined as the difference between the number of end users assigned to the instance with the highest and lowest amount of end users, respectively. Several statistical properties of these objective functions are presented in Table II. Additionally, Figure 1 displays the cumulative distribution function of objective values that are attained across all placements.

Due to the fact that the latency measures are continuous, they yield significantly more distinct values, resulting in smooth CDF curves. In contrast, the imbalance is always an integer value which is constrained by the number of nodes in the topology. Hence, individual steps are visible in the plot. Since the average latency between end users and services is calculated from 34 individual latencies, outliers are smoothed out and the resulting variance is relatively low. The values of the maximum latency objective have a higher variance and fewer distinct values since the maximum does not necessarily change between similar placements that share multiple controller locations.

TABLE II: Various statistics of the objective functions.

Objective	Number of distinct values	Mean	Variance
$\pi^{\text{avg latency}}$	45,311	0.195	0.001
$\pi^{\text{max latency}}$	244	0.491	0.013
$\pi^{\text{imbalance}}$	29	0.305	0.019

IV. WEIGHTING METHODS

In order to aggregate the performance of a placement that is evaluated with respect to multiple objective functions into a single value, the mechanisms that are analyzed in this work require weights for each considered dimension. Hence, we first discuss methods for obtaining these weights based on the set of placements and the corresponding objective values.

In the following, the weight of the j -th objective is denoted as w_j and weights are normalized, i.e., $\sum_{j=1}^m w_j = 1$ in case of m objective functions. Additionally, objective values are also normalized prior to applying the weighting mechanisms. The observed values for n placements and m objective dimensions are stored in an $n \times m$ matrix A which is transformed into the normalized matrix R according to Equation 1.




Fig. 1: Empirical CDFs of objective values.

$$r_{ij} = \frac{a_j^{max} + a_j^{min} - a_{ij}}{a_j^{max} + a_j^{min}} \quad (1)$$

In this equation, $a_j^{min} = \min_i a_{ij}$ and $a_j^{max} = \max_i a_{ij}$ refer to the minimum and maximum values of the j -th objective, respectively.

A. Uniform Weighting

As a baseline naïve approach, we use a weighting mechanism that does not take into account any observed data and assigns equal weights to every objective, i.e., $w_j^{uni} = \frac{1}{m}$.

B. Entropy-Based Weighting

In information theory, (the Shannon) entropy is used as a means to quantify the amount of information that is stored in a message [21]. The key idea behind the entropy-based weighting method consists of assigning higher weights to objective dimensions that carry more information, i.e., those that have a higher number of distinct values and low individual occurrence probabilities for each value. Based on [22], the weights are calculated in three steps. First, observed values are normalized for each dimension (cf., Equation 2).

$$p_{ij} = \frac{r_{ij}}{\sum_{i=1}^n r_{ij}}, \quad j \in \{1, \dots, m\} \quad (2)$$

Then, the entropy is determined by means of

$$e_j = -\frac{1}{\ln n} \sum_{i=1}^n p_{ij} \ln p_{ij}, \quad j \in \{1, \dots, m\}. \quad (3)$$

Finally, the weight is calculated as

$$w_j^{ent} = \frac{1 - e_j}{\sum_{i=1}^m (1 - e_i)}, \quad j \in \{1, \dots, m\}. \quad (4)$$

C. Weighting Based on the Coefficient of Variation

Intuitively, objectives whose values cover a wide range of different values tend to have a higher impact on the total resulting performance of a placement than objectives that attain only few values or values that are very close to each other. Hence, we investigate the suitability of the coefficient of variation for quantifying the relative importance of an objective. The coefficient of variation is defined as the ratio between the standard deviation and the mean of observed values. Thus, the weights are calculated according to Equation 5. σ_j and μ_j refer to the standard deviation and mean of the j -th objective, respectively.

$$w_j^{cv} = \frac{\frac{\sigma_j}{\mu_j}}{\sum_{i=1}^m \frac{\sigma_i}{\mu_i}}, \quad j \in \{1, \dots, m\} \quad (5)$$

D. Weighting Based on the Standard Deviation

Similarly to the weighting approach that is based on the coefficient of variation, this mechanism uses the standard deviation in order to calculate the relative weights.

$$w_j^{sd} = \frac{\sigma_j}{\sum_{i=1}^m \sigma_i}, \quad j \in \{1, \dots, m\} \quad (6)$$

E. Comparison

In order to allow for a comparison between the different weighting mechanisms, Figure 2 presents the weights of individual objectives according to the four weighting approaches. The x-axis denotes the objective and is ordered according to Table II, i.e., dimension 1 corresponds to $\pi^{\text{avg latency}}$. The height and color of the bars represent the weight and weighting method, respectively.

While the weights that are returned by the different mechanisms differ in terms of absolute values, the relative order of objectives is consistent. Having the lowest variance and the narrowest interquartile range, the latency between end users and services is assigned the lowest weights. As discussed in Section III, the maximum-based measure has a higher variance and thus also results in higher weights when compared to its average-based counterpart. The highest weights are assigned to the imbalance measure. This can be explained by the high variance that is observed in the context of the imbalance objective.

A comparison of the absolute weights that are assigned by the weighting methods shows that the mechanisms that are based on standard deviation and the coefficient of variation return similar values. This phenomenon can be explained by the fact that objective values are normalized prior to applying the weighting methods. Thus, the normalization using the mean that is applied in the context of the latter does not have a large impact on the final weights. Finally, the entropy-based weighting approach yields the widest range of weights, i.e., between less than 0.1 and more than 0.6. This indicates a higher sensitivity towards the objectives' variance, which seems to be the main influence factor on the resulting weight for all weighting methods that take into account observed objective values.




Fig. 2: Relative weights of objectives according to different weighting mechanisms.

V. RANKING METHODS AND RESULTS

A. Ranking Methods

To aggregate the scores a_{ij} of the different attributes j of the placement i to an overall ranking score ρ_i , four well-known multi-attribute decision methods will be considered.

First, we consider Simple Additive Weighting (SAW) [23], which computes the overall score by adding the normalized attribute scores $r_{ij} = \frac{a_j^{min}}{a_{ij}}$ multiplied by the weights w_j .

$$\rho_i^{SAW} = \sum_j w_j \cdot r_{ij}$$

A similar ranking method is Multiplicative Exponent Weighting (MEW) [24], which calculates the overall score as the product of the normalized attribute scores $r_{ij} = \frac{a_j^{min}}{a_{ij}}$, which are given the respective weight as exponent.

$$\rho_i^{MEW} = \prod_j r_{ij}^{w_j}$$

The Technique for Order Preference by Similarity to Ideal Solution (TOPSIS) [23] normalizes the attributes $r_{ij} = \frac{a_{ij}}{\sum_i a_{ij}}$, and computes the distances to an optimal placement with all best weighted normalized attribute values $v_j^{min} = \min_i(w_j r_{ij})$, and to a worst placement composed of all worst weighted normalized attribute values $v_j^{max} = \max_i(w_j r_{ij})$. Then, the separation between the optimal and the worst placement is computed by

$$s_i^{min} = \sqrt{\sum_j (w_j r_{ij} - v_j^{min})^2}$$

$$s_i^{max} = \sqrt{\sum_j (w_j r_{ij} - v_j^{max})^2}$$

The resulting ranking ρ_i is the relative closeness to the ideal solution:

$$\rho_i^{TOPSIS} = \frac{s_i^{max}}{s_i^{min} + s_i^{max}}$$

VIKOR [25] relies on the best and worst attribute values, a_j^{min} and a_j^{max} . Then, for each placement, scores are calculated by two strategies:

$$S_i = \sum_j w_j \frac{a_j^{min} - a_{ij}}{a_j^{min} - a_j^{max}}, \quad R_i = \max_j \left(w_j \frac{a_j^{min} - a_{ij}}{a_j^{min} - a_j^{max}} \right)$$

The final ranking score for each placement is then computed with a parameter γ , $0 \leq \gamma \leq 1$, for the weight of each strategy, and the best and worst values of S_i and R_i , i.e., $S^{min} = \min_i S_i$, $S^{max} = \max_i S_i$, $R^{min} = \min_i R_i$, $R^{max} = \max_i R_i$:

$$\rho_i^{VIKOR} = \gamma \frac{S_i - S^{min}}{S^{max} - S^{min}} + (1 - \gamma) \frac{R_i - R^{min}}{R^{max} - R^{min}}$$

We set $\gamma = 0.5$ to give equal weight to both strategies.

Together with the four weighting methods presented in Section IV, this gives 16 different ranking methods for the multi-objective placement problem. Due to the vast amount of distinct placements, we will apply the 16 methods only to the subset of Pareto-optimal placements, i.e., the set of placements in which no attribute can outperform any other attribute

B. Comparison of Resulting Rankings

TABLE III: Gordon α .

Method 1	Method 2	α
(ρ^{SAW}, w^{ent})	(ρ^{MEW}, w^{ent})	10
(ρ^{SAW}, w^{ent})	(ρ^{TOPSIS}, w^{ent})	10
(ρ^{SAW}, w^{ent})	(ρ^{VIKOR}, w^{ent})	10
(ρ^{MEW}, w^{sd})	(ρ^{TOPSIS}, w^{sd})	10
(ρ^{MEW}, w^{ent})	(ρ^{VIKOR}, w^{ent})	10
(ρ^{MEW}, w^{uni})	(ρ^{TOPSIS}, w^{sd})	10
(ρ^{TOPSIS}, w^{ent})	(ρ^{VIKOR}, w^{ent})	10
(ρ^{SAW}, w^{ent})	(ρ^{VIKOR}, w^{uni})	4
(ρ^{SAW}, w^{uni})	(ρ^{VIKOR}, w^{sd})	4
(ρ^{MEW}, w^{sd})	(ρ^{VIKOR}, w^{uni})	4
(ρ^{MEW}, w^{uni})	(ρ^{VIKOR}, w^{uni})	4
(ρ^{TOPSIS}, w^{sd})	(ρ^{VIKOR}, w^{uni})	4
(ρ^{TOPSIS}, w^{uni})	(ρ^{VIKOR}, w^{uni})	4

Table IV lists the highest and lowest correlations between different combinations of weighting and ranking methods in terms of Kendall's τ and Spearman's ρ rank order correlation coefficients. It can be seen that generally high correlations can be achieved between all ranking algorithms. In contrast, small negative correlation can be seen only for VIKOR with uniform weights. Thus, this might give some evidence that the investigated algorithms mainly agree on the inherent order of the elements.

Another metric for measuring the agreement between rankings was proposed by Gordon [12]. Gordon's α is defined as the number of objects, which are contributing to the agreement

TABLE IV: Highest and lowest correlations between different combinations of weighting and ranking methods.

Method 1	Method 2	τ	ρ
(ρ^{SAW}, w^{ent})	(ρ^{MEW}, w^{ent})	1.00	1.00
(ρ^{SAW}, w^{ent})	(ρ^{TOPSIS}, w^{ent})	1.00	1.00
(ρ^{SAW}, w^{ent})	(ρ^{VIKOR}, w^{ent})	1.00	1.00
(ρ^{MEW}, w^{sd})	(ρ^{TOPSIS}, w^{sd})	1.00	1.00
(ρ^{MEW}, w^{ent})	(ρ^{TOPSIS}, w^{ent})	1.00	1.00
(ρ^{MEW}, w^{ent})	(ρ^{VIKOR}, w^{ent})	1.00	1.00
(ρ^{MEW}, w^{uni})	(ρ^{TOPSIS}, w^{uni})	1.00	1.00
(ρ^{TOPSIS}, w^{ent})	(ρ^{VIKOR}, w^{ent})	1.00	1.00
(ρ^{SAW}, w^{sd})	(ρ^{VIKOR}, w^{uni})	-0.11	-0.16
(ρ^{SAW}, w^{ent})	(ρ^{VIKOR}, w^{uni})	-0.11	-0.15
(ρ^{MEW}, w^{ent})	(ρ^{VIKOR}, w^{uni})	-0.11	-0.15
(ρ^{TOPSIS}, w^{ent})	(ρ^{VIKOR}, w^{uni})	-0.11	-0.15
(ρ^{VIKOR}, w^{ent})	(ρ^{VIKOR}, w^{uni})	-0.11	-0.15

between the rankings: $\alpha := N - \delta$. Thus, it can be computed as the difference between the length of the ranking N and the minimum number of objects δ , which have to be removed to ensure a perfect agreement between the reduced rankings. Gordon's α confirms the high correlation coefficients, as there are many pairs of rankings with a perfect agreement of $\alpha = N = 10$. The lowest value of α is 4, which shows that still the ranking order is not completely inverted by any algorithm-weighting combination.

Probabilistic ranking models give another approach to comparing the obtained rankings. Luce [14] constructs probabilities for a ranking $\rho = (i_1, i_2, \dots, i_N)$ from conditional probabilities. Thus, after $r - 1$ stages, p_{i_r} is defined as the probability that the element i_r is the most preferred element from the set of remaining elements $B = \{i_r, \dots, i_N\}$. By repeating the choice, this gives the probability of the rating ρ as:

$$P(\rho) = \prod_{r=1}^{N-1} \frac{p_{i_r}}{\sum_{j \in B} p_j}$$

The highest Luce probabilities are obtained by a ranking, which was created by the combinations $(\rho^{TOPSIS}, w^{sd}|w^{cv}|w^{uni})$ and $(\rho^{MEW}, w^{sd}|w^{uni})$. This means, this ranking gives high ranks to the elements, which are most preferred by all algorithm-weighting combinations. Note that this ranking is also the modal ranking in the resulting set of rankings. All four entropy based algorithms output the same ranking, which reaches the second highest Luce probabilities. Towards the other end, the SAW and VIKOR algorithms and the the standard deviation (sd) and coefficient of variation weighting (cv) output rankings with low probabilities (with the above mentioned exceptions).

Mallow's Φ -model is based on paired comparison of the ranked elements. It can be formulated as

$$P_{\rho_0, \theta}(\rho) = \left(\sum_{\rho} \theta^{X(\rho_0, \rho)} \right)^{-1} \cdot \theta^{X(\rho_0, \rho)}, \quad 0 \leq \theta < \infty,$$

in which $X(\rho_0, \rho)$ is Kendall's τ distance, i.e., the number of disagreements between ρ_0 and ρ . ρ_0 is an a priori set location parameter (e.g., the modal ranking or an averaged ranking), and θ is a measure of variation, which will be fitted from the rankings with a table given in [16]. Following the methodology presented by Feigin and Cohen in [16], the model also allows to detect outlier rankings. Using the averaged ranking as location parameter and fitting θ accordingly, the highest probability is obtained by the ranking of (ρ^{MEW}, w^{cv}) . The second highest probabilities are achieved by the modal ranking, which already accounted for the highest Luce probabilities. Again the entropy rankings have the third highest probability. This means that these three rankings are closest to the averaged ranking, which was chosen as location parameter. Using the modal ranking as location parameter, the order of the first and second rating would change, but the entropy rating would still receive the third highest probability. The outlier detection, which mainly depends on the fitting of θ , indicates that (ρ^{MEW}, w^{cv}) is an outlier ranking with a too high probability, and (ρ^{SAW}, w^{uni}) and (ρ^{VIKOR}, w^{uni}) are outliers with a too low probability close to 0. In particular, this means that the disagreements for (ρ^{SAW}, w^{uni}) and (ρ^{VIKOR}, w^{uni}) are exceptionally high compared to the averaged or modal ranking.




Fig. 3: Pareto optimal placements and their ranks according to the presented ranking mechanisms.

Following the approach described in [17], Figure 3 shows a boxplot of the ranks of the different placements sorted by median. It can be seen that there are small boxes for the first five placements, which means that there is a large agreement among the different algorithm-weighting combinations. Only for the last five placements, there is some disagreement among the different rankings. Still several outliers can be observed, however, taking a detailed look at the data, most outlier ratings stem from uniform weighting of the attributes. Thus, this weighting method seems to be inappropriate for ranking the placements.

To sum up, the different ranking methods showed a high agreement, especially for the top-ranked placements.

This means, among the investigated methods, no algorithm-weighting stands out and most of them are well suited to combine the Pareto-optimal placements into a single score. Nevertheless, the results suggest that the use of uniform weights can lead to outlier rankings, which do not reproduce the majority rankings.

VI. CONCLUSION

In this work, we applied multi-objective decision methods to the problem of selecting the best placement for a cloud service from a set of Pareto-optimal placements. Therefore, we investigated four methods to determine the relative importance of different objectives (i.e., uniform, entropy-based, coefficient of variation-based, and standard deviation-based weighting), and four multi-objective optimization algorithm (i.e., simple additive weighting, multiplicative exponent weighting, TOPSIS, and VIKOR). We showed that, for the investigated problem, most algorithm-weighting combinations perform sufficiently good and have a high level of agreement, especially on the top-ranked placements. Only the usage of uniform weights was shown to cause outlier rankings, which, nevertheless, can provide a complementary view on the ranked placements. In case of all approaches, the weights and ranks for a given set of placements can be efficiently calculated by implementing the presented equations. However, operators might need to employ heuristic approaches for large problem instances in order to find feasible placements in a timely manner.

In future work, we will investigate other multi-objective optimization algorithms and other placement problems, e.g., on other network topologies, or with other/additional attributes. Thereby, also the computational complexity has to be taken into account, as some algorithms need an impractically high runtime for the ranking of a large numbers of placements. Eventually, the goal will be to derive guidelines for choosing the appropriate ranking mechanism for Pareto-optimal placements in the cloud service placement problem.

ACKNOWLEDGMENTS

This work was partly funded in the framework of the EU ICT project INPUT (H2020-2014-ICT-644672) and the CELTIC EUREKA project SENDATE-PLANETS (Project ID C2015/3-1), which is partly funded by the German BMBF (Project ID 16KIS0474). The authors alone are responsible for the content.

REFERENCES

- [1] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The Internet Topology Zoo," *IEEE JSAC*, vol. 29, no. 9, 2011.
- [2] B. Jennings and R. Stadler, "Resource Management in Clouds: Survey and Research Challenges," *Journal of Network and Systems Management*, vol. 23, no. 3, pp. 567–619, 2015.
- [3] M. Bari, S. R. Chowdhury, R. Ahmed, R. Boutaba *et al.*, "On Orchestrating Virtual Network Functions," in *11th International Conference on Network and Service Management (CNSM)*, Barcelona, Spain, 2015, pp. 50–56.
- [4] M. C. Luizelli, L. R. Bays, L. S. Buriol, M. P. Barcellos, and L. P. Gaspary, "Piecing Together the NFV Provisioning Puzzle: Efficient Placement and Chaining of Virtual Network Functions," in *IFIP/IEEE International Symposium on Integrated Network Management (IM)*, Ottawa, ON, Canada, 2015, pp. 98–106.
- [5] H. Moens and F. De Turck, "VNF-P: A Model for Efficient Placement of Virtualized Network Functions," in *10th International Conference on Network and Service Management (CNSM)*, Rio de Janeiro, Brazil, 2014, pp. 418–423.
- [6] S. Mehraghdam, M. Keller, and H. Karl, "Specifying and Placing Chains of Virtual Network Functions," in *3rd International Conference on Cloud Networking (CloudNet)*, Luxembourg, Luxembourg, 2014, pp. 7–13.
- [7] M. F. Bari, A. R. Roy, S. R. Chowdhury, Q. Zhang, M. F. Zhani, R. Ahmed, and R. Boutaba, "Dynamic Controller Provisioning in Software Defined Networks," in *9th International Conference on Network and Service Management (CNSM)*, Zurich, Switzerland, 2013, pp. 18–25.
- [8] R. Guerzoni, R. Trivisonno, I. Vaishnavi, Z. Despotovic, A. Hecker, S. Beker, and D. Soldani, "A Novel Approach to Virtual Networks Embedding for SDN Management and Orchestration," in *Network Operations and Management Symposium (NOMS)*, Krakow, Poland, 2014, pp. 1–7.
- [9] S. Lange, S. Gebert, T. Zinner, P. Tran-Gia, D. Hock, M. Jarschel, and M. Hoffmann, "Heuristic Approaches to the Controller Placement Problem in Large Scale SDN Networks," *IEEE Transactions on Network and Service Management - Special Issue on Efficient Management of SDN and NFV-based Systems*, vol. 12, 2015.
- [10] S. Lange, S. Gebert, J. Spoerhase, P. Rygielski, T. Zinner, S. Kounev, and P. Tran-Gia, "Specialized Heuristics for the Controller Placement Problem in Large Scale SDN Networks," in *International Teletraffic Congress (ITC 27)*, Ghent, Belgium, 2015.
- [11] M. G. Kendall, *Rank correlation methods*. Griffin, 1948.
- [12] A. Gordon, "A measure of the agreement between rankings," *Biometrika*, vol. 66, no. 1, pp. 7–15, 1979.
- [13] L. L. Thurstone, "A law of comparative judgment," *Psychological review*, vol. 34, no. 4, p. 273, 1927.
- [14] R. D. Luce, *Individual choice behavior*. John Wiley, 1959.
- [15] C. L. Mallows, "Non-null ranking models. i," *Biometrika*, vol. 44, no. 1/2, pp. 114–130, 1957.
- [16] P. D. Feigin and A. Cohen, "On a model for concordance between judges," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 203–213, 1978.
- [17] A. Cohen, "Analysis of large sets of ranking data," *Communications in Statistics-Theory and Methods*, vol. 11, no. 3, pp. 235–256, 1982.
- [18] A. Cohen and C. L. Mallows, "Assessing goodness of fit of ranking models to data," *The Statistician*, pp. 361–374, 1983.
- [19] P. Diaconis, "Group representations in probability and statistics," *Lecture Notes-Monograph Series*, vol. 11, pp. i–192, 1988.
- [20] D. E. Critchlow, M. A. Fligner, and J. S. Verducci, "Probability models on rankings," *Journal of mathematical psychology*, vol. 35, no. 3, pp. 294–318, 1991.
- [21] C. E. Shannon, "A Mathematical Theory of Communication," *The Bell System Technical Journal*, 1948.
- [22] J. Huang, "Combining entropy weight and TOPSIS method for information system selection," in *IEEE Conference on Cybernetics and Intelligent Systems*, 2008.
- [23] C.-L. Hwang and K. Yoon, *Multiple Attribute Decision Making - Methods and Applications: A State-of-the-art Survey*. Springer Science & Business Media, 1981, vol. 186.
- [24] S. H. Zanakis, A. Solomon, N. Wishart, and S. Dublisch, "Multi-attribute Decision Making: A Simulation Comparison of Select Methods," *European Journal of Operational Research*, vol. 107, no. 3, pp. 507–529, 1998.
- [25] S. Opricovic and G.-H. Tzeng, "Compromise Solution by MCDM Methods: A Comparative Analysis of VIKOR and TOPSIS," *European Journal of Operational Research*, vol. 156, no. 2, pp. 445–455, 2004.