Mobile Edge Computing: Challenges for Future Virtual Network Embedding Algorithms

Michael Till Beck, Marco Maier
Ludwig-Maximilians-Universität München
München, Germany
{michael.beck,marco.maier}@ifi.lmu.de

Abstract—**Mobile edge computing aims at reducing network latency and network stress by deploying mobile applications at the network edge. This paper proposes network virtualization in the context of mobile edge computing networks as an enabler for future, flexible and shared network infrastructures (IaaS). Network virtualization leads to the Virtual Network Embedding (VNE) problem, which aims at deploying virtual networks onto a shared, physical infrastructure. This paper is a position paper discussing new challenges for future VNE algorithms. To this end, new parameters specific to the mobile edge computing scenario are analyzed which are not considered by state-of-the-art VNE algorithms. Furthermore, novel and edge-specific VNE optimization objectives are derived.**
Keywords—Virtual Network Embedding, Next-gen Cellular Networks, Edge Computing

I. INTRODUCTION

Starting with the introduction of the first smartphones, one of the most apparent challenges for mobile network operators is handling future bandwidth demands, which are expected to further increase dramatically over the next years [1].

Mobile data traffic is predicted to continue doubling each year. Video contributes heavily to overall mobile network traffic. The usage of mobile applications (apps) accessing services deployed in the Internet is expected to further contribute to this trend, resulting in a growth of around 12 times by 2018 [2]. This trend becomes even more remarkable as novel mobile devices (like Google Glass and other, wearable devices) and applications arise. Offering more and more hardware capabilities, these devices pave the way for novel application types like augmented reality [2].

Network operators spend enormous efforts to keep up with these demands in order to satisfy the needs of their users, providing low-latency network access. Upgrading base stations and core network routers to higher-capacity equipment reduces high utilization in the core network, but comes with significant operational cost. Furthermore, network operators are impelled to quickly integrate upcoming technologies (like Long Term Evolution, LTE) at the edge of their networks, offering better quality of experience. Higher bandwidth capacities at the network edge, however, directly affect core network utilization and require further investments.

Two technologies have recently been proposed as a remedy to this dilemma: 1) Mobile edge computing and 2) network virtualization. Mobile edge computing aims at reducing latency by shifting computational efforts from the Internet cloud to the mobile edge. Network virtualization increases management flexibility of the mobile infrastructure and enables resource sharing between multiple providers.

1) *Mobile edge computing (MEC)* is an emerging technology that is seen as an alleviating factor in this context [3], [4], [5]. MEC aims at reducing both network latency and resource demands by shifting computing and storage capacities from the Internet cloud to the mobile edge. Instead of uploading or downloading content generated or demanded by the mobile user, mobile applications refer to a service located close to the current position of the user. Services are hosted on devices directly attached to base stations or smart cells (i.e., macrocells, microcells, or picocells). These hosts are also known as MEC servers and are operated by the mobile infrastructure provider. The proximity of the MEC servers to the mobile device not only takes load from the mobile core but also increases responsiveness of mobile applications.

2) *Network virtualization* is commonly seen as a key technology for the Future Internet and has recently been proposed in the context of mobile networks [1], [6], [7], [8]. Network Virtualization enables sharing of physical network resources like base stations, core routers, and MEC servers between multiple network operators. Network virtualization enables operators to fully isolate their (virtual) resources from those hosted by others on the same physical device (data and control plane). Network sharing not only reduces cost for deployment of new hardware resources but also operational cost [1][6][8]. Increasing network management flexibilities, virtualization of network resources is also seen as a key technology to mitigate the ossification of the core protocols.

This paper proposes a fully virtualized MEC infrastructure. Virtualizing both the mobile core and the mobile edge network enables infrastructure providers to shared resources between several mobile operators (Infrastructure as a service, IaaS), including computing and storage capacities of the MEC servers servers, and enhances management flexibilities. One major challenge of network virtualization is the embedding of virtualized resources onto the physical network. This problem is known as the VNE problem. VNE algorithms aim to embed multiple Virtual Network Requests (VNRs) onto a shared substrate network, enabling virtual network operators to share a common substrate infrastructure flawlessly by assigning sufficient resources. While virtualization is a technique that is well-known both in mobile networks and computer networks, the VNE problem has not been discussed in the context of MEC so far. Therefore, this paper addresses this shortcoming and identifies new research directions for future, MEC-specific VNE approaches.

The remainder of this paper is structured as follows: In section II, network topologies of MEC networks are explained. Furthermore, the VNE problem is formalized and the formalization is extended with respect to the mobile edge scenario. Section III-A introduces new VNE parameters and Section III-B introduces VNE optimization objectives for MEC networks. In section III-C, challenges for future VNE approaches are discussed. Section IV discusses related work and section V concludes the paper.

II. VIRTUALIZATION OF MOBILE EDGE INFRASTRUCTURES

This section discusses mobile edge computing, motivates network virtualization in the context of mobile networks, and

 ISBN: 978-1-61208-354-4

formulates the VNE problem in the context of mobile edge networks.

A. Mobile Edge Computing

Mobile edge computing is an emerging concept becoming more and more feasible with the shift towards the LTE wireless communication standard. LTE's new core network, *System Architecture Evolution (SAE)*, is an all-IP network with a simplified architecture, allowing for greater flexibility of the network's topology and more heterogeneous access networks, integrating legacy systems (e.g., air interfaces of GPRS or UMTS) and LTE's new *Evolved Universal Terrestrial Radio Access (E-UTRA)*. Due to LTE's low-latency and high-bandwidth radio access networks, deployment of new computing resources at the mobile edge becomes a promising approach for supporting novel latency-sensitive applications.

MEC servers provide computing, storage and bandwidth capacity that is shared by multiple virtual machines installed on top of them. Fig. 1 depicts the mobile edge computing scenario. MEC servers, being owned and managed by the infrastructure provider, are directly attached to the base stations. Traditionally, all data traffic originating at the data centers is forwarded by Internet routers to the mobile core network. The traffic is routed through the core network to a base station which delivers the content to the mobile devices. In the mobile edge computing scenario, MEC servers take over some or even all of the tasks originally performed in a data centers. Being located at the mobile edge, this eliminated the need of routing these data through the core network, leading to low communication latency.

Two different, but related usage scenarios have been proposed in the context of mobile edge computing: The first one proposes mobile devices to delegate calculations to the MEC servers (*offloading* of resource- or power-intense tasks) [5], while the second one proposes application service providers (ASPs) to deploy services traditionally hosted within data centers on the MEC servers (*Edge Deployment*) [3], [4]. Both aim at making the edge of the mobile network smarter, leading to a reduction of core network utilization and decreased latency:

1) Offloading Some applications running on a mobile device are capable of offloading resource- or power-intense tasks to MEC servers. Therefore, the mobile application invokes additional services deployed at a virtual machine hosted on a MEC servers. MEC servers are placed nearby, offer excellent Internet connectivity, and are easily reachable by the mobile device: In the best case, there is a one-hop communication between mobile device and host, offering low-latency access. The concept is expected to increase limited computing, storage, or bandwidth capacities of mobile devices by referring to external, resource-rich resources. Another objective of offloading is to reduce power consumption of a mobile device.

If no MEC server is available, the mobile device degrades gracefully to a more distant MEC server, a remote Internet cloud server, or use its own hardware resources [5], [9].

2) Edge Deployment In the edge deployment scenario, network providers offer MEC capacities for the deployment of ASP-operated virtual machines running at the edge. ASPs offer additional services at the network edge, increasing responsiveness of their applications. Services or parts of services traditionally hosted in data centers are now shifted to the network edge. Since traffic between MEC servers and mobile devices has not to be routed through the core network, this leads to decreasing core network utilization and lower communication latency.

B. Network Virtualization

Network virtualization has been proposed both in the context of computer networks and for mobile core networks [3], [4], [5]. This paper proposes the application of network virtualization techniques for the whole network infrastructure, including network core, base stations, and MEC servers. Network virtualization is proposed as a key technology to overcome the ossification of core protocols, since it enables the deployment of several, isolated virtual networks on top of a shared physical infrastructure. Virtual networks are co-hosted on a common substrate infrastructure and, since they are fully isolated, are even capable of deploying different communication protocols (e.g., IPv4/v6 or proprietary protocols) on the same substrate links.

Infrastructure providers (InPs) offer physical network resources to several mobile network operators. Operators specify network topologies and hardware resource demands to be deployed within the infrastructure of the InP. Operators are usually external customers of the InP. This does not exclude, however, that InP itself can also deploy networks on its own on top of its infrastructure, renting spare resources to other operators. The InP provides its physical resources to the operators, ensuring that all network requests of the operators are fulfilled. In this paper, fully-virtualized MEC networks are proposed. This means that both the network core, and also the network edge, i.e., MEC servers and base stations resources provide virtualization capabilities. Network virtualization is a useful technique in order to separate several internal networks and to increase manageability. Furthermore, it enables the InP to rent spare resources to other operators.

One important aspect in this area of research lies in the embedding of virtual network entities to the physical (or to be more general: the substrate) infrastructure. This is commonly known as the Virtual Network Embedding (VNE) problem [10]. Physical resources are limited and have to be shared between the virtual network entities that are assigned to these resources. This is depicted in Fig. 2: two network requests are assigned to a substrate network. The VNE problem is divided in two sub-problems: *Virtual Node Mapping* and *Virtual Link Mapping*: Virtual nodes are assigned to substrate nodes offering sufficient resources. Virtual links are either assigned to a single substrate link, or span a path of multiple links in the substrate network, where each link offers sufficient resources. This is shown in Fig. 2 for the virtual link demanding 100MBit/s bandwidth capacity. The VNE problem becomes $\mathcal{NP}$-hard when substrate nodes and links have finite resources [10].

C. Problem Formulation

In this subsection, a formal description for the general VNE problem as depicted in Fig. 2 is presented. This formal model will then be enhanced with respect to MEC specific properties.

A substrate network $S = (N, L)$ is modeled as a set of substrate nodes N and a set of links L mutually connecting some of the nodes. Similar to the substrate network, a VNR is modeled as a collection of virtual nodes N^i and links L^i. Substrate nodes and links offer resources R, assigned by $\mathrm{cap} : N \cup L \rightarrow 2^R$. Virtual nodes and links *demand* these resources, formally described as $\mathrm{dem}_i : N^i \cup L^i \rightarrow 2^R$. The objective of a VNE algorithm is to embed several Virtual Network Requests VNRs, denoting $\mathrm{VNR}^i = (N^i, L^i)$ as being the i-th request. Virtual entities that are embedded onto a substrate entity *consume* substrate resources they demanded. Therefore, the VNE has to assure that a sufficient amount of resources is provided by a substrate entity before a virtual entity gets assigned to it. The embedding is modeled as a

ISBN: 978-1-61208-354-4

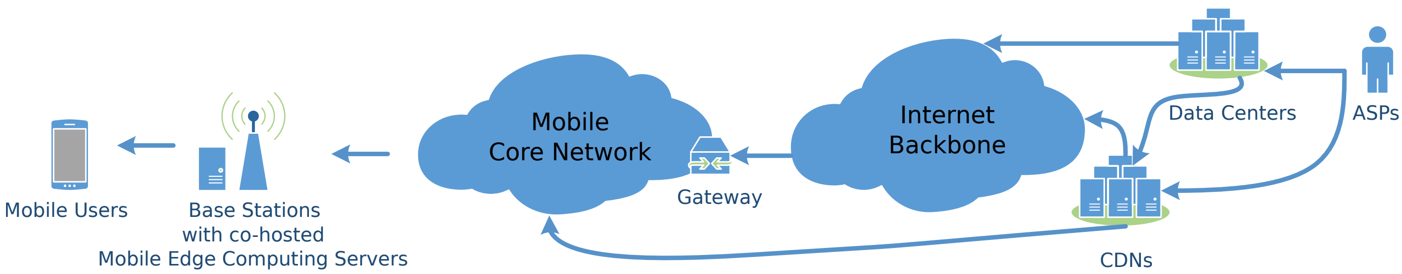

Fig. 1: Mobile Edge Computing: Deployment of MEC servers at the Edge of the Mobile Network

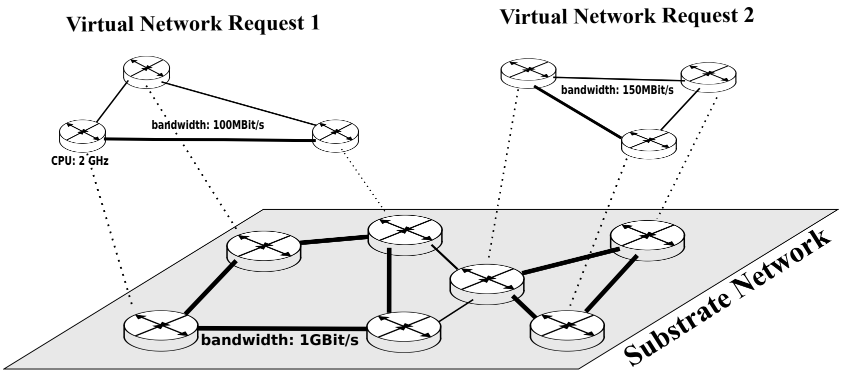

Fig. 2: Virtual Network Embedding

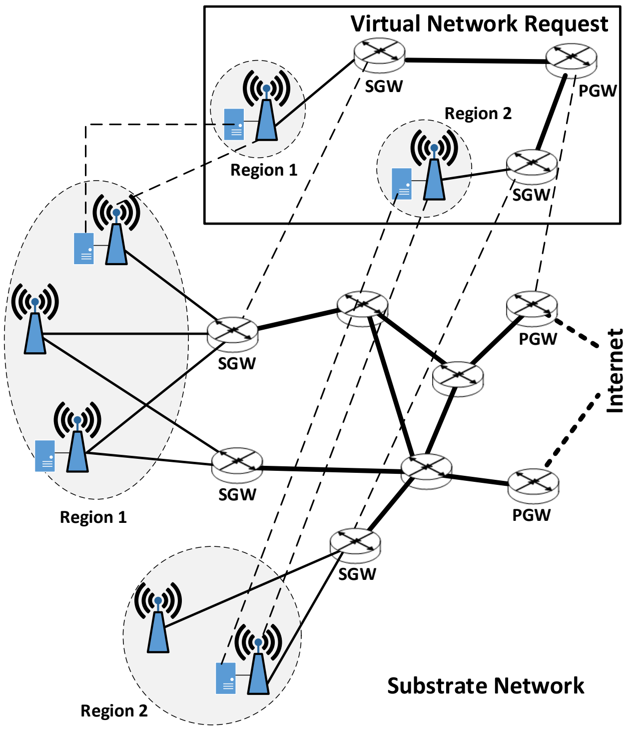

Fig. 3: Embedding of Mobile Edge Computing Networks

function $f_i : N^i \rightarrow N$ assigning nodes of VNR^i to the substrate network, and a function $g_i : L^i \rightarrow \mathrm{SN}' \subseteq SN$ assigning edges of the VNR to the substrate links (or a combination of links, i.e., paths). The notation used here is in line with the one introduced in [10].

The VNE problem in the mobile edge scenario is depicted in Fig. 3: For this scenario, a differentiation between edge and core nodes is needed. Therefore, the following new variables are introduced: N_{edge} and N_{core} for substrate edge/core nodes, L_{edge} and L_{core} for substrate edge/core links, and, respectively, N^i_{edge}, N^i_{core}, L^i_{edge}, and L^i_{core} for the virtual nodes and edges.

Edge nodes are nodes located at the edge of the network, i.e., nodes providing network access to the mobile devices (representing base stations) or computing and storage capabilities (representing MEC servers). Core nodes are all other nodes (network core devices), e.g., routers, *Serving Gateways (SGW)* (connecting the base stations to the core network) or *Packet Data Network Gateways (PGW)* (connecting the core network to the internet). A main characteristic of edge nodes representing base stations is their relatedness to a specific geographic region. Fig. 3 depicts the embedding of a VNR. In this case, the VNR demands base stations with affiliated MEC servers in two specific regions (region 1 and 2), thus not allowing to, e.g., only use the two base stations with MEC capabilities in region 1. Geographic constraints of edge nodes are one of the new challenges for VNE in the MEC scenario, and are reflected in the proposed VNE parameters and optimization objectives which are discussed in the next section.

III. Embedding Virtual Networks in Mobile Edge Computing Scenarios

This section introduces the VNE problem in the context of mobile edge computing. To this end, the VNE problem (1) is extended with respect to edge-specific parameters and (2) new optimization objectives are introduced. Furthermore, (3) general challenges for future VNE algorithms in mobile edge networks are discussed. To the best of our knowledge, these parameters, optimization objectives, and mobile edge related challenges have not been discussed in literature before.

A. *New VNE parameters*

VNE algorithms take different parameters of network resources into account for calculating network embeddings. As an example for such parameters, substrate resources have individual capacities, and virtual resources have respective requirements (also known as *demands*). Various kinds of resources are

Copyright (c) IARIA, 2014. ISBN: 978-1-61208-354-4

assigned to nodes and links: For example, in traditional VNE scenarios, CPU resources are assigned to nodes and bandwidth resources to links. Such a strict distinction in node and link parameters, however, is problematic, since some parameters influence each other. For example, CPU utilization of a node being part of a path between two nodes influences available bandwidth on that communication path.

Fischer et al. propose a classification of VNE parameters: *Primary/secondary* and *functional/non-functional* [10]. Primary parameters like CPU resources are directly assigned to a substrate resource. Secondary parameters, however, depend on other primary parameters and these side-effects have to be considered first. As an example, packet loss at a node depends on the primary parameters CPU resource and memory resource. Functional parameters like CPU-/bandwidth capacity specify low-level functionality. Non-functional parameters are high-level properties. For example, resilience and security are both non-functional parameters.

Being two well-known parameters common to most existing VNE approaches, CPU and bandwidth also apply in the context of MEC. Both core and edge nodes provide CPU capacities. On core nodes, CPU capacity is consumed for tasks like routing and mobility-specific tasks like authorization and billing. Edge nodes, i.e., MEC servers, provide CPU capacity for applications deployed by the ASPs. In the following, new VNE parameters are introduced that are applicative to the MEC scenario. Parameters are classified as being primary/secondary and functional/non-functional.

MEC Coverage (primary & functional) Base stations provide wireless link capacity, connecting mobile devices to the mobile network. The range of a typical macrocell base station covers several dozens of kilometers. Other types of base stations, small cells like microcells, cover much lower ranges. Relay nodes are base stations providing enhanced coverage at cell edges and hot-spot regions, covering the same region as the main base station.
Coverage refers to the geographical region covered by the base stations of the same mobile operator, i.e., the geographical region in which mobile devices are able to communicate with the operator's core network. Coverage depends on lower level characteristics such as physical location of the base station, transmission power and environmental influences. Increasing the transmission power of a base station increases the radius of the covered region, but also decreases average bandwidth available to the mobile users. Increasing coverage leads to interferences with other base stations if frequencies are reused. In the MEC scenario, coverage also influences availability of resources provided by MEC servers due to the fact that MEC servers are directly connected to base stations. Thus, increasing coverage also increases availability of MEC resources. However, since more mobile users are able to reach the MEC server, this leads to higher utilization of the MEC server and, thus, to higher latency. Therefore, there is a tradeoff between availability and utilization/latency of MEC resources.

MEC Server Storage (primary & functional) MEC servers provide disk storage for virtual machines running on top of them in order to, e.g., cache proximity-related data. This requirement is not considered by most existing VNE approaches.

Latency (secondary & functional) If the VNE algorithm embeds a latency-sensitive link to a path spanning several substrate links, the sum of all latency properties of these links may never exceed the demanded maximum latency of the corresponding virtual link. Latency is a critical factor for several mobile applications. Therefore, operators define an upper bound for reaching a MEC server from a certain base station. Depending on this value, the VNE is either able to choose between any nearby MEC server or is forced to use substrate resources of a MEC server right at the given base station. Communication latency is a secondary, functional parameter, since it depends on CPU utilization and refers to low-level functionality.

Regional Bandwidth Capacity (secondary & functional) Cell capacity refers to available bandwidth resources provided by one or more base stations within a region (up- and downlink). Bandwidth capacity is increased by operating multiple *geographically adjacent* base stations, each covering a subset of the same region, instead of just one macrocell. Furthermore, LTE Advanced allows the parallel deployment of small cells within the *same* region that is covered by a macrocell (enhanced Inter-cell interference coordination, eICIC), leading to improved bandwidth capacity. In the MEC scenario, regional bandwidth capacity is proportional to the amount of available MEC server resources that are accessible within a geographical region. Bandwidth capacity, being a secondary parameter, is limited by bandwidth resources of deployed base stations and by bandwidth resources of the links connecting the base stations to the core network.

Regional MEC Computing and Storage Capacity (secondary & functional) Similar to regional bandwidth capacity, CPU/storage capacity refers to CPU/storage resources available in a region. CPU/storage capacity compounds of the resources of MEC servers directly connected to the base stations deployed in that region and computing capacity of *logically adjacent* MEC servers, i.e., MEC servers that are either directly connected to base stations covering that region or well-connected through the mobile network to the base stations.

MEC Resources per Mobile Device (secondary & functional) Since all mobile devices in a single cell share resources corresponding to the above parameters CPU capacity, storage, bandwidth and latency, VNRs demand resources on a per-user basis, e.g., a certain minimum amount of bandwidth per user in order to achieve the intended quality of experience. As a consequence, the VNE has to lower the average number of users in certain cells, which is done by adjusting the coverage of adjacent cells (i.e., to distribute the users among multiple smaller cells). Vice versa, when resource requirements are smaller than available resources, the VNE is able to utilize a smaller number of base stations, omitting certain nodes completely, allowing to (temporarily) shut down these nodes for energy-saving purposes.

The next subsection discusses MEC-specific optimization objectives for the VNE problem. Due to the fact that many MEC-related parameters are secondary, i.e., they depend on other, primary parameters –, the interdependencies are emphasized.

B. New Optimization Objectives

VNE algorithms compute the *optimal* (or near-optimal) embedding of a set of VNRs. An embedding is optimized with respect to an optimization objective. As an example, such an optimization criterion is, in order to reduce cost for the IsP, to minimize the number of substrate resources that are utilized by the embedded virtual networks. This section discusses various VNE optimization objectives in the context of MEC.

Increase MEC Coverage One main objective for network operators is to provide cell coverage to their mobile users. Operators both aim at expanding geographical dimension of cell coverage, and the capacity of their cells: The more bandwidth resources are available in a region, the more users are able to connect to their networks (and pay for using their

services). Therefore, one objective for future VNE algorithms is to share resources provided by the base stations optimally between VNRs with respect to coverage and bandwidth. However, this is not limited to base stations: In MEC, this also applies to resources provided by MEC servers. Cells providing bandwidth resources also provide computing capabilites. These capabilities have to be adequately shared between network operators covering the same regions.
There is a tradeoff between power consumption and coverage, since more resources are needed to cover larger areas or to provide higher bandwidth capacities.

Reduce Latency Reducing communication latency is of major importance for today's mobile network operators. VNE algorithms are obliged to take this aspect into account by embedding VNRs in a way that minimizes latency at the edge and/or in the core network. Reducing latency at the edge, i.e., between MEC servers and base stations, leads to small response times between services hosted at the MEC servers and mobile users. To reduce overall network delay, delay between base stations and gateway nodes connecting the network core to the Internet (PGWs, cf. Fig. 3) has to be considered by future VNE approaches.

Provide QoS-compliant Embeddings Besides latency, future VNE algorithms have to consider flexible QoS- and QoE-considerations. VNRs either demand strict resource assignments which have to be reserved exclusively. As an example, a VNR requests 1km cell capacity in a specific area and 10GB of MEC server storage capacity. These resources are then exclusively reserved for the VNR and thus can not be shared with other VNRs. As an alternative, VNE algorithms should balance utilization throughout the substrate network resources in order to guarantee equivalent QoS for all VNRs. In this case, VNRs do not request such strict assignments and the VNE aims at providing well-balanced cell coverage and MEC server capacities for all VNRs.

Maximize Regional MEC Capacities Regional bandwidth/ computing/ storage capacity refers to available resources provided by one or more base stations or MEC servers within a region. Bandwidth is limited by resources of the base stations, but also by the links connecting the base station to the core network. Bandwidth capacity is increased by overlapping coverage of *geographically adjacent* base stations. Computing capacity is increased by utilizing *logically adjacent* MEC servers, i.e., MEC servers that are well-connected through the core network to a base station covering the region.

Provide Resilience / Fault Tolerance for MEC services Resilience at the network edge is provided by a VNE by allocating resources of additional, geographically adjacent base stations covering the same region. If another base station fails, this backup resource takes over. The VNE has to ensure that the backup resource does not introduce any interferences with other base stations. Backup MEC server capabilities are provided by allocating logically adjacent MEC server resources. The path from the base station to the backup MEC server should provide similar latency conditions as the old one.

Reduction of Power Consumption Only few VNE approaches aim at reducing power consumption of substrate resources [10]. However, since they focus on data center routers and servers, these approaches are only partly applicable in the context of MEC. In the MEC scenario, new power models and interdependencies have to be considered: Increasing transmission power of base stations extends coverage, but also influence other base stations [11]. Decreasing power consumption decreases coverage, impelling mobile devices to re-register to other base stations due to poor connection quality to the current BS. This leads to load shifting both at the edge of the mobile network and, as a consequence, also in the network core. Current energy-aware VNE approaches aim at switching off as many substrate entities as possible in order to decrease overall power consumption. However, in the MEC scenario, this not only results in decreased coverage, but also in reduction of MEC server computing and storage capacity. If the connection of the base station to the core network is shut down, the MEC server gets isolated. MEC servers connected to base stations that are already in use should be preferred instead of switching on MEC servers connected to inactive core resources.
Therefore, there is a tradeoff that should be taken into account by future VNE algorithms: Coverage of base stations, available bandwidth, and MEC server capacity in this region vs. energy efficiency. How sparse can network infrastructure be / how many cloudlets can be switched off, while still ensuring enough bandwidth and CPU capacities?

Provide Security at the Edge Due to security considerations, not all operators or ASPs want to share MEC servers running their applications with other operators. Therefore, virtual MEC servers of different providers are embedded onto different substrate MEC servers in order to reduce the risk of any malicious influence (e.g., denial of service attacks). The challenge for new VNE approaches is to find an embedding that considers security constraints by still ensuring latency/QoS demands.

C. Challenges in Mobile Edge Networks

This subsection outlines several general challenges for VNE algorithms in the context of large-scale, fully-virtualized mobile networks.

Scalability As mentioned before, the VNE problem is $\mathcal{NP}$-hard. Therefore, most VNE approaches are based on heuristics. This results in non-optimal solutions, but decreases the size of the problem and leads to significant improvements with respect to runtime. However, almost all algorithms rely on a single, central node calculating the VNE. Centralization hinders scalability for large-scale, dynamic networks. In fact, most algorithms were evaluated with substrate networks spanning only few dozens or hundreds of nodes. Such settings might be realistic in the context of middle sized networks like testbeds. However, it is far away from national-wide or transnational mobile networks. One might be puzzled by the fact, therefore, why only few approaches are distributed: In fact, almost all VNE algorithms are centralized and require full knowledge of the substrate network topology and substrate resources [10], [12]. In large-scale environments like real-life mobile networks, current VNE algorithms are stretched to their limits [13].

Mobility-Awareness of Embeddings Another VNE challenge in MEC networks is to provide mobility-awareness. Mobile users move between neighbored cells, resulting in additional handoff overhead and routing of traffic through different paths. Therefore, embedding communication paths of virtual base stations covering adjacent cells on similar core network paths reduces variations in latency (jitter). A mobility-aware VNE algorithm aims at providing both physical and logical proximity for virtual base stations covering adjacent regions.

Utilization-Awareness of Embeddings In mobile networks, both edge and core utilization fluctuates significantly depending on the time of day. Demand also varies due to big sports events, new year's Eve, etc. Davy et al. propose the usage of user mobility models for shared network resources, predicting aggregated movement patterns of mobile users. Furthermore, popularity prediction methods for video content discussed, predict distribution of video content based on their expected popularity [13]. Integrating such prediction methods is seen as promising step towards utilization-awareness VNE algorithms. Being static and centralized, most VNE algorithms cannot cope with dynamic and flexible VNR requirements in large-scale environments. Therefore, novel VNE algorithms should be both dynamic and distributed in order to provide flexibility and adaptability in large network scenarios.

IV. Related Work

This section discusses related work: Several papers related to mobile network virtualization are mentioned. To the best of our knowledge, there is neither related work on network virtualization in the context of edge computing networks nor on VNE algorithms in this context.

Virtualization-based isolation techniques in the context of mobile networks are proposed as an enabling technology for future, cost-efficient mobile networks, shared and operated by multiple network operators. Different network infrastructure sharing scenarios are discussed and sharing options are classified based on different business models [1].

A general virtualization-enabled network architecture is proposed in [6]. The advantages of shared, hereogeneous network infrastructures are emphasized and applications of network virtualization in this context is discussed. An approach towards the virtualization of LTE networks has been introduced by [8]. Authors discuss the advantages of virtualizing LTE infrastructures and elaborate on virtualization of the LTE air interface.

An extensive and up-to-date classification of current VNE parameters and objectives is given in [10]. Many VNE approaches have been proposed so far, focussing on embedding objectives like cost-optimization, resilience etc. While some objectives are related to the ones presented in this paper (e.g., resilience and security), none considers the special demands of MEC networks. Since InPs and network operators usually aim for a combination of multiple (possibly contrary) objectives (e.g., reducing power consumption to a certain extend by also providing a sufficient degree of network resilience), an in-depth evaluation of these approaches in combination with the novel VNE parameters and objectives presented in this paper is left for future work.

V. Conclusion

This paper proposes network virtualization as a key technology for future mobile edge computing networks. The Virtual Network Embedding problem is analyzed in this context and new challenges for future embedding algorithms are discussed. To the best of our knowledge, no VNE approach has been published so far considering the MEC-specific network parameters and optimization objectives presented here. Therefore, the authors hope that this position paper provides an initial step towards new VNE approaches. As a first step, the authors are currently implementing a delay-aware VNE algorithm that considers geographical constraints as well as latency considerations.

References

[1] A. Khan, W. Kellerer, K. Kozu, and M. Yabusaki, "Network sharing in the next mobile network: TCO reduction, management flexibility, and operational independence," Communications Magazine, IEEE, vol. 49, no. 10, pp. 134–142, 2011.

[2] Ericcson, "Ericcson mobility report june 2013," http://www.ericsson.com/res/docs/2013/ericsson-mobility-report-june-2013.pdf [retrieved: 2014-01-28].

[3] Intel and Nokia Siemens Networks, "Increasing mobile operators' value proposition with edge computing," http://nsn.com/system/files/document/edgecomputingtechbrief_328909_002_0.pdf [retrieved: 2014-01-28].

[4] IBM Corporation, "Smarter wireless networks; add intelligence to the mobile network edge," http://public.dhe.ibm.com/common/ssi/ecm/en/wsw14201usen/WSW14201USEN.PDF [retrieved: 2014-01-28].

[5] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The case for vm-based cloudlets in mobile computing," Pervasive Computing, IEEE, vol. 8, no. 4, pp. 14–23, 2009.

[6] M. Hoffmann and M. Staufer, "Network virtualization for future mobile networks: General architecture and applications," in Communications Workshops (ICC), 2011 IEEE International Conference on. IEEE, 2011, p. 1–5.

[7] E. Kudoh and F. Adachi, "Power and frequency efficient virtual cellular network," in Vehicular Technology Conference, 2003. VTC 2003-Spring. The 57th IEEE Semiannual, vol. 4. IEEE, 2003, pp. 2485–2489. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1208838

[8] Y. Zaki, L. Zhao, C. Goerg, and A. Timm-Giel, "LTE mobile network virtualization: Exploiting multiplexing and multi-user diversity gain," Mobile Networks and Applications, vol. 16, no. 4, pp. 424–432, 2011.

[9] D. Fesehaye, Y. Gao, K. Nahrstedt, and G. Wang, "Impact of cloudlets on interactive mobile cloud applications," 16th International Enterprise Distributed Object Computing Conference, pp. 123–132, Sep. 2012.

[10] A. Fischer, J. F. Botero, M. T. Beck, H. De Meer, and X. Hesselbach, "Virtual network embedding: A survey," Communications Surveys & Tutorials, IEEE, vol. 15, no. 4, pp. 1888–1906, 2013.

[11] F. Richter, A. J. Fehske, and G. P. Fettweis, "Energy efficiency aspects of base station deployment strategies for cellular networks," in Vehicular Technology Conference Fall. IEEE, 2009, p. 1–5.

[12] M. T. Beck, J. F. Botero, A. Fischer, H. De Meer, and X. Hesselbach, "A distributed, parallel, and generic virtual network embedding framework," in IEEE Int'l Conf. on Communications (ICC). IEEE, 2013. [Online]. Available: http://www.net.fim.uni-passau.de/pdf/Beck2013a.pdf

[13] S. Davy, J. Famaey, J. Serrat, J. L. Gorricho, A. Miron, M. Dramitinos, P. M. Neves, S. Latre, and E. Goshen, "Challenges to support edge-as-a-service," Communications Magazine, IEEE, vol. 52, no. 1, pp. 132–139, 2014.