

A Prioritised Traffic Embedding Mechanism enabling a Public Safety Virtual Operator

Jonathan van de Belt*, Hamed Ahmadi*, Linda E. Doyle*, Oriol Sallent†

*The Centre for Future Networks and Communications - CONNECT, Trinity College Dublin

†Universitat Politècnica de Catalunya

Abstract—Public Protection and Disaster Relief (PPDR) services can benefit greatly from the availability of mobile broadband communications in disaster and emergency scenarios. While undoubtedly offering full control and reliability, dedicated networks for PPDR have resulted in high operating costs and a lack of innovation in comparison to the commercial domain. Driven by the many benefits of broadband communications, PPDR operators are increasingly interested in adopting mainstream commercial technologies such as Long Term Evolution (LTE) in favour of expensive, dedicated narrow-band networks.

In addition, the emergence of virtualization for wireless networks offers a new model for sharing infrastructure between several operators in a flexible and customizable manner. In this context, we propose a virtual Public Safety (PS) operator that relies on shared infrastructure of commercial LTE networks to deliver services to its users. We compare several methods of allocating spectrum resources between virtual operators at peak times and examine how this influences differing traffic services. We show that it is possible to provide services to the PS users reliably during both normal and emergency operation, and examine the impact on the commercial operators.

I. INTRODUCTION

Mobile broadband communications have great potential to improve the efficiency of Public Protection and Disaster Relief operations by enabling applications such as real-time access to high-resolution maps and floor plans, on-field live video transmission from helmet cameras to a central unit, telemedicine, etc. [1] [2]. Driven by this demand, significant changes are taking place in the professional mobile radio industry. While existing Professional/Private Mobile Radio (PMR) technologies, such as TETRA, TETRAPOL, DMR, APCO P25, GSM-R and others, have been very successful in delivering business and mission critical voice and narrowband data services in these professional sectors, these technologies are not well suited to support higher data rate applications. In this context, the adoption of mainstream commercial technologies such as LTE/LTE-Advanced for PPDR mission critical communications is gaining strong momentum [3]. Indeed, many national and international PPDR organizations have endorsed or are considering LTE as the next generation technology either to augment their existing systems, or to provide a future migration path. The economies of scale of LTE are expected to bring down the costs of private mobile radio equipment. Establishing common technical standards for the PMR and commercial domains offers significant opportunities for exploiting the synergies between these two domains through network and spectrum sharing principles [4]. Though LTE

technology was originally conceived and designed as a flexible and spectrally efficient mobile broadband technology for the commercial domain, work is in progress within the 3GPP to extend the use of LTE to critical communication applications and to deliver future services to PPDR users [5].

In most parts of the world, the prevailing model that has been used for the delivery of mission critical narrowband PPDR communications (voice-centric and low data rate services) is based on the use of dedicated networks; networks deployed for PPDR use that are mainly privately built and operated with the specific purpose of serving the communications needs of a limited number of agencies. This delivery model, while undoubtedly offering full control and high availability of communications resources to PPDR users, has resulted in a niche market with far less innovation and higher prices for communication equipment in comparison to the commercial domain.

Given that PPDR communications systems are fundamentally funded from constrained public authorities' budgets (i.e., taxpayer's money), PPDR communications should strive to achieve the most cost-effective solution. Without compromising the high control, security and resilience standards required by PPDR communications, future broadband PPDR data services must be delivered through optimal deployment mechanisms, in a cost-effective manner with a sustainable business model [5].

Acknowledging that while the most suitable model may differ across countries and regions because of differing existing conditions and interests, this paper considers the delivery of mobile broadband services to PPDR users over public-access LTE-based networks. These networks are run by commercial operators and the PPDR users share the infrastructure with the rest of commercial traffic.

Despite some reluctance to use commercial networks from the PPDR sector, since it is believed that current commercial networks cannot satisfy many of their requirements (see e.g., [6] [7]), there is increasing consensus that these infrastructures undoubtedly have a role to play in critical communications solutions. As commercial broadband networks become an important part of society's infrastructure, they can enable PPDR users to experience the benefits of broadband communications (see e.g., [8] [9]).

In terms of materialization of these principles, [4] presents a system architecture that enables PPDR service access across a number of LTE-based dedicated and commercial network.

The system is implemented through a mobile virtual network operator (MVNO) model such that no dedicated LTE networks are deployed for PPDR. As an example of practical realization, ASTRID, the national operator of radio communications, paging and the dispatching network for emergency and security services in Belgium, launched a broadband data service called Blue Light Mobile [10] in Spring 2014. This service allows PPDR users to use the commercial 3G networks for data-centric applications. To that end, ASTRID takes on the role of an MVNO that offers services via the mobile networks of the three national operators in Belgium.

In turn, sustained on the emergence of virtualization techniques, [11] proposes a shared physical infrastructure where multiple heterogeneous virtual operators (VOs) coexist concurrently in a much more sophisticated and resource-efficient manner than is the case for current MVNOs. In this model, VOs are able to request resources in real-time from a physical infrastructure provider (InP) based on detailed requirements, and the virtual embedding problem is dynamically addressed through embedding algorithms that can re-embed existing virtual networks to achieve more efficient resource allocation and satisfy additional resource requests.

Given these emerging capabilities, this paper investigates whether a virtual PS operator can exist on a virtualized commercial LTE network. We examine the reliability of such a network for PS services, and the benefits to the network owner.

II. SYSTEM OVERVIEW

A. Wireless Network Virtualization

In wireless network virtualization full isolation is required between different virtual networks, so that each virtual network has the illusion of full control over its resources. Virtual networks should not be able to interfere with each other. Link isolation can be achieved by dividing the physical resources into orthogonal dimensions to prevent interference [12].

For this work we consider the virtualization of a single enhanced Node B (eNodeB) in a LTE network. In this model, a hypervisor is added on top of the physical hardware of the eNodeB that abstracts the physical eNodeB resources and manages the resources for several virtual eNodeBs [13]. The hypervisor is also responsible for assigning the LTE spectrum resources (known as physical resource blocks (PRBs)) amongst the virtual networks. Since the PRBs consist of two orthogonal dimensions, frequency and time, full isolation can be ensured between virtual networks.

On this physical eNodeB several VOs can exist, each with their own virtual eNodeBs. For now we assume that sufficient computational resources exist to satisfy the needs of the virtual eNodeBs, and that the computational resources are fully isolated. We are interested only in the PRB assignment.

At regular time intervals or rounds, the hypervisor performs the PRB assignment. We assume that the time between consecutive rounds is such that each of the PRBs have the same average throughput and are equally adequate in the long run.

This means the hypervisor can assign any of the PRBs to a virtual network rather than assigning specific resource blocks.

At any time, VOs can make virtual resource requests (VRRs) to the hypervisor on behalf of their users for r PRBs. To ensure that this model applies to both uplink and downlink resource allocation in LTE, the PRBs resources allocated to each user must be contiguous [14]. Therefore, upon receiving a virtual resource request for r resources, the hypervisor maps this to a rectangle $f \times t$, with time and frequency dimensions f and t respectively, which ensures that the physical resource blocks assigned to each user are contiguous.

In addition to specifying the number of resources desired, VRRs also specify a duration, d , for the number of rounds that the resources are desired. A priority level, p , is also assigned to VRRs, which allows VOs to prioritise requests depending on the service type, s . Thus each request for a service s , $VRR(s)$, has three parameters: $\{p, r, d\}$.

Arriving VRRs are stored in a buffer until the next resource allocation round. At each round, an attempt is made to allocate resources to all virtual resource requests in the buffer, within the constraints of the total resources available, $F \times T$, where F and T are the frequency and time dimensions of the substrate. This process is also known as the embedding process. If there are enough free resources to embed a VRR, the request is embedded and deemed successful. The virtual operator that made the request can then use the resources to satisfy a user service for the specified number of rounds (d). If a VRR is not successfully embedded during a round, it remains in the buffer and is embedded in the next round if possible. A VRR that is not embedded for a number of rounds, max_delay , is removed from the buffer and the VRR is rejected. max_delay can vary for different services and priority levels.

B. Embedding Algorithms

We examine several embedding algorithms that the hypervisor could employ to maximize the resource occupancy of the substrate and minimize the number of VRRs that are rejected. The resource allocation problem is an NP-complete knapsack problem with the additional complexity of priority levels. A complete mathematical formulation of the embedding problem with priority levels is given in our previous work [11].

We investigate static and dynamic embedding algorithms based on the heuristic Karnaugh-map algorithm proposed in [15]. In this algorithm VRRs are embedded onto a two-dimensional substrate by finding Karnaugh-map like regions of vacant resources. The algorithm attempts to cluster the VRRs together, so that the number of contiguous unused resources is as large as possible, which allows additional virtual resource requests of varying sizes to be embedded. At each round, the set of VRRs in the buffer are ordered by priority and decreasing area. Then for each VRR, the Karnaugh map approach is used to find a list of free resources with dimensions equal to or greater than the dimensions of the VRR. Next the smallest region that can contain the request is chosen. However, within this region it is necessary to find the best corner at which to embed the VRR.

The Embedding Density Index (EDI) is used to find the corner that maximizes resource clustering. The EDI measures the number of edges between free and occupied resource blocks. A high EDI reflects a substrate where the occupied resources are spread out, whereas a substrate in which the occupied resources are grouped together has a low EDI. The effect of embedding the VRR at a corner is tested by calculating the resulting EDI. This is done for each of the four corners in the region chosen to contain the request and the corner with the lowest EDI is selected as the embedding location. This ensures that the new virtual service is clustered among the existing virtual services to the greatest degree possible. Further details are provided in [15].

In the static case, successful VRRs receive the same set of resources every round. However with this static embedding scheme it is possible that a virtual network request is rejected due to topology limitations even though there are enough resources available to satisfy it. To overcome this drawback, we consider the dynamic version of the Karnaugh map algorithm, from our previous work [11]. It dynamically reassigns the existing services every round. The resources can be used much more efficiently and additional VRRs can be accepted.

C. Priorities during Normal and Emergency operation

The hypervisor has two modes of operation: normal operation and emergency operation. During normal operation a virtual resource request from a given service VRR(s) has an associated priority level p . When an emergency occurs, the priority levels can be changes on a per-service and per VO basis. Thus, for example, during normal operation the Public Safety VO may be treated in the same manner as any other commercial operator, but during emergencies, certain services of the Public Safety VO can be given higher priority.

III. INVESTIGATED SCENARIO

The scenario we are investigating is that of a commercial LTE-based network, where the physical infrastructure provider has virtualized the mobile resources and virtual operators can request these resources as needed. In our case there is one virtual PS network operator and one virtual commercial VO that share the physical network.

A. Evaluation Metrics

The metrics of interest in this scenario are the percentage of virtual resource requests that are rejected per service and the percentage of total resources used per service at every round. The rejection rate for a round is calculated as the sum of $r \times d$ of rejected VRRs, divided by the sum of $r \times d$ of all VRRs. The first metric shows the importance of reliability from the perspective of the VOs, while the latter metric provides the InP with a method of measuring how efficient its resources are being used. Based on these metrics the InP could decide whether additional investment in physical infrastructure is required, and virtual operators can discover whether the InP provides a satisfactory service level.

B. Traffic Modelling

We model three common types of services that each operator desires to fulfil for its users, namely voice traffic, video traffic, and messaging traffic, corresponding to 3GPP's CQI classes 2, 3, and 6 respectively [16, Table 6.1.7]. The number of resources, r , required for a service is modelled as a uniform distribution $U(min, max)$, and this is then mapped to a permutation of the smallest rectangular area covering at least r PRBs. The duration of each VRR, d , follows an exponential distribution with an average lifespan of μ rounds.

We model the three services as follows:

$$\begin{aligned} \text{VRR}(\text{voice}) &= \{p, d = 30, r \sim U(min = 1, max = 2)\} \\ \text{VRR}(\text{video}) &= \{p, d = 10, r \sim U(min = 8, max = 25)\} \\ \text{VRR}(\text{msg}) &= \{p, d = 3, r \sim U(min = 1, max = 8)\} \end{aligned}$$

The maximum delay, max_delay , is set to 1 round for voice service requests, 2 rounds for video services, and 4 rounds for messaging services. The values for these parameters reflect realistic conditions to the best of our knowledge.

C. Service Priorities

During normal operation, the VRRs made by the PS and commercial VOs are considered as having equal importance and priority levels are based only on the service type: the highest priority is for voice requests, next is video and then messaging. When an emergency occurs, all PS traffic is given higher priority than commercial traffic, retaining the same service priority levels as before. Regarding commercial traffic, messaging traffic is considered more important during an emergency than video traffic. To this end, higher priority is given to voice, then messaging and eventually video.

D. Simulation setup

In this scenario the dimensions of the substrate are set to $F = 20, T = 20$. The resource embedding is performed for 1000 rounds. Initially, normal, day-to-day operation is assumed for the hypervisor, however after 300 rounds an emergency occurs. This means that the priority levels of VRRs change as described above. This emergency situation lasts for 400 rounds, after which the hypervisor switches back into normal operation.

The number of VRRs of each service arriving per round is modelled as a Poisson process with the following average rates λ requests per round: For normal operation, it takes the values 1.4, 1.4, and 3 for commercial voice, video and messaging respectively, and the values 0.14, 0.14, and 0.3 for PS voice, video and messaging. These values offer a realistic approach, since the total percentage of each service type is similar to what is found in the literature and in real systems today [17]. Regarding the traffic share between commercial and PS, it assumes that PS represents 10% of the traffic. When an emergency occurs, the number of VRRs increases significantly. The PS now requires many more resources and we multiply the values by 5 for all PS services. The commercial voice and messaging traffic are also assumed to increase by a factor of 2.5, while we consider that the video traffic remains constant.

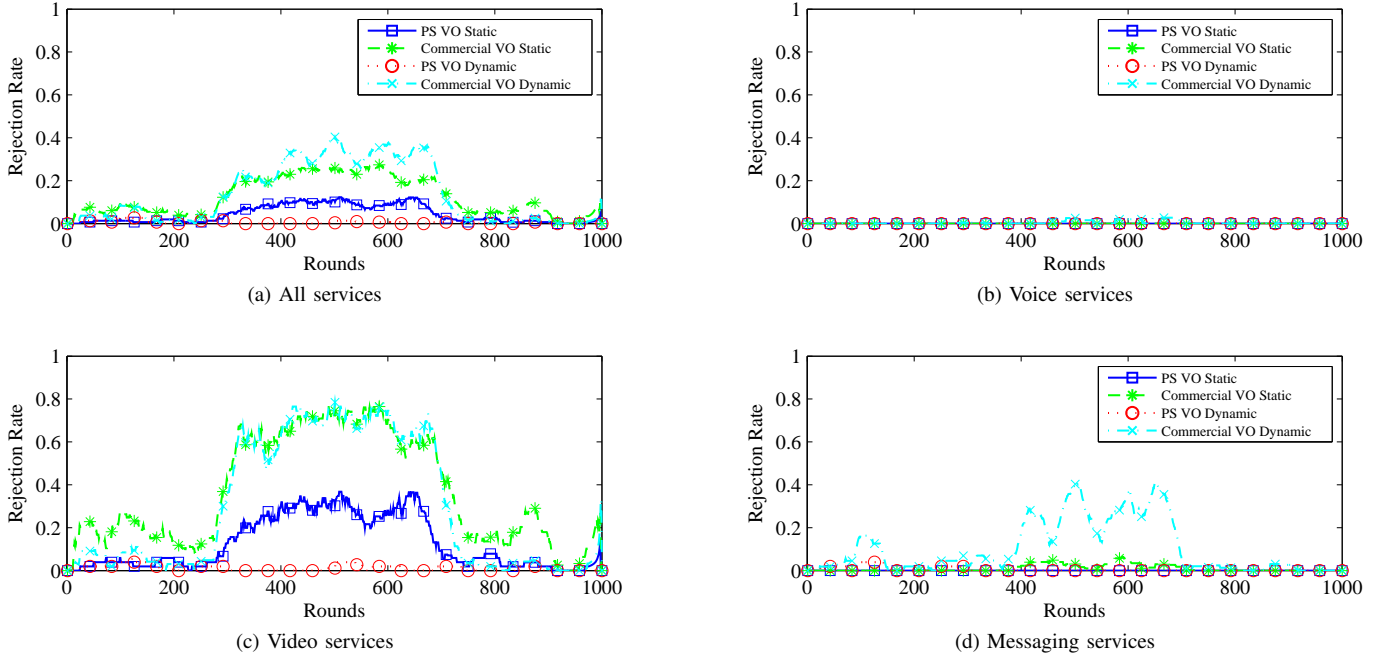


Fig. 1. Even during an emergency (rounds 300-700), negligible rejection rates for the Public Safety services can be achieved for all services (a) using the dynamic embedding algorithm, while for the commercial operator voice services (b) are completely reliable, messaging services (d) are quite reliable, but video services (c) are unreliable.

IV. PERFORMANCE RESULTS

We compare the performance of the embedding algorithms in normal and emergency operation and examine whether virtualized LTE networks can offer adequate service for PS.

The performance achieved in terms of rejection rate per round is depicted in Figure 1. Regarding the performance of the embedding algorithms, it can be observed in (a) that the dynamic embedding achieves 0% rejection for all the offered PS traffic during the emergency, while the static algorithm rejects about 10% of the PS VRRs. Since lower priority has been assigned to the commercial traffic during the emergency, higher rejection rates (in the order of 30% - globally across all commercial services) are observed. Based on these rejection rates, PS communication can be achieved reliably over this system.

In more detail, regarding the treatment of the different service flows during the emergency, (b) reflects that all voice traffic from both PS and commercial is served. For video service, (c) reveals that the lowest priority given to commercial video traffic during the emergency leads to high rejection rates, in the order of 70%. Regarding PS video traffic, it is observed that the dynamic embedding algorithm is able to provide almost 0% rejection rate, while the static embedding leads to substantial rejections (above 20%). Finally, for the messaging service, (d) shows that rejections only occur for commercial traffic and the dynamic embedding algorithm. This is because the dynamic embedding has led to increased allocation of resources to PS traffic.

The performance achieved in terms of resources used per

round is depicted in Figure 2. The comparison of (a) and (b) illustrates how the dynamic algorithm is able to exploit almost 100% of the resources during the emergency, while the static embedding exhibits lower resource exploitation. During normal operation, the resource usage does not reach 100% because the offered traffic during this period is below the maximum networks capacity. It can also be observed how the capacity share for the PS during the emergency increases (it represents around 40% of the traffic served during the emergency) compared to the normal operation. This is because of the PS traffic increase during the emergency period (i.e., higher requested load) as well as the higher priority assigned to PS traffic (i.e., higher served load). The distribution of the networks capacity across services is shown in (c) and (d). Both the static and dynamic embedding algorithm devote almost 50% of the capacity to voice during the emergency. However, the dynamic embedding is able to allocate a substantially larger share of video traffic (about 25% of the total traffic) compared to the static approach, at the expense of a slight reduction of messaging traffic. In any case, as shown in (d), messaging traffic for the dynamic embedding benefits from the ability of the algorithm to reach almost 100% resource occupation during the emergency period.

V. CONCLUSION

In this work we proposed a virtual public safety operator that shares a commercial LTE network with virtual commercial operators. Several embedding algorithms were analysed in terms of the rejection rate and the percentage of resources occupied. We examined the reliability of such a system in

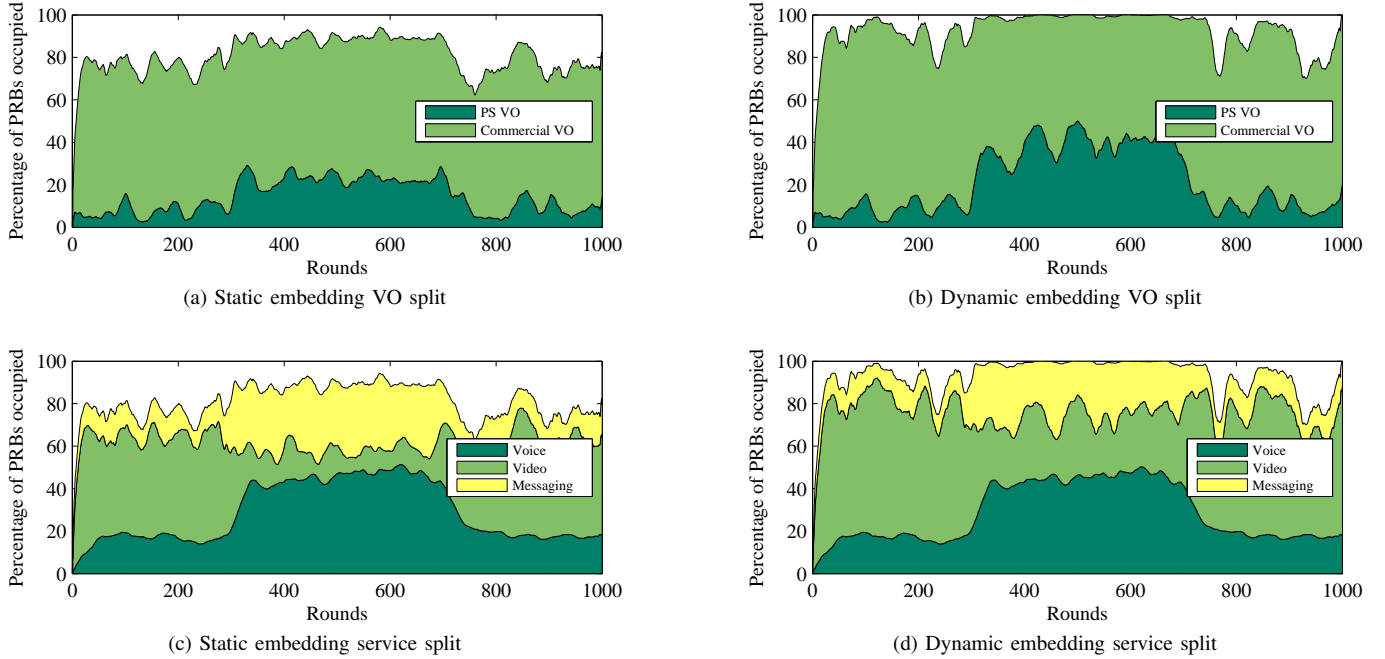


Fig. 2. Compared to the static embedding (a), the Public Safety operator receives a significantly higher share of total resources when dynamic embedding is used (b). This is mostly due to the low number of video services that are embedding in the static case (c), in comparison to the dynamic case (d).

normal and emergency situations, and how differing traffic services were affected. We showed that it is possible to provide the public safety operator with reliable communication during emergencies, since almost zero services were rejected. The impact on commercial services was also investigated, and though the performance dropped during emergencies, the most crucial services (voice and messaging) were very reliable.

These promising results motivate further studies, such as the formulation of utility functions to capture users acceptance of the provided service levels during normal and emergency conditions. This would allow more complete design and parametrization of the embedding algorithm. Additionally, the proposed public safety virtual operator model can be detailed further according to actual/future LTE architectures and 3GPP-based prioritization and Quality-of-Service control mechanisms.

ACKNOWLEDGMENTS

This work is supported by Science Foundation Ireland under Grants No. 10/CE/I1853, No. 13/RC/2077, and 10/IN.1/I3007 and the Spanish Research Council and FEDER funds under ARCO grant (ref. TEC2010-15198). Special thanks to Dr. Ismael Gomez and Dr. Paul Sutton.

REFERENCES

- [1] U.S. NPSTC Public Safety Communications Report, "Public Safety Communications Assessment 2012-2022," Tech. Rep., 2012.
- [2] 2CEPT, "User requirements and spectrum needs for future European broadband PPDR systems (Wide Area Networks)," ECC, Tech. Rep., 2013.
- [3] TCCA Critical Communications Broadband Group, "Mission Critical Mobile Broadband : Practical standardisation & roadmap considerations," Tech. Rep. 4155039, 2013.
- [4] R. Ferrus, O. Sallent, G. Baldini, and L. Goratti, "Lte: the technology driver for future public safety communications," *Communications Magazine*, IEEE, vol. 51, no. 10, pp. 154–161, October 2013.
- [5] R. Ferrus and O. Sallent, "Extending the lte/lte-a business case: Mission and business-critical mobile broadband communications," *Vehicular Technology Magazine*, IEEE, vol. 9, no. 3, pp. 47–55, Sept 2014.
- [6] TCCA Critical Communications Broadband Group, "PPDR need for dedicated communications and why they cannot rely on public systems for mission critical communications," Tech. Rep., 2013.
- [7] ETSI, "Additional spectrum requirements for future Public Safety and Security (PSS) wireless communication systems in the UHF frequency range," Tech. Rep., 2010.
- [8] TCCA, "Mobile Broadband for Critical Communications Users: A review of options for delivering Mission Critical solutions," Tech. Rep. 4155039, 2013.
- [9] Erikson, "Public safety mobile broadband," Tech. Rep., 2014.
- [10] Blue Light Mobile Service, <http://bluelightmobile.be/en/home>.
- [11] J. van de Belt, H. Ahmadi, and L. Doyle, "A dynamic embedding algorithm for wireless network virtualization," in *Vehicular Technology Conference (VTC Fall)*, 2014 IEEE 80th, Sept 2014, pp. 1–6.
- [12] K.-m. Park and C.-k. Kim, "A framework for virtual network embedding in wireless networks," *Proceedings of the 4th International Conference on Future Internet Technologies - CFI '09*, p. 5, 2009.
- [13] Y. Zaki, L. Zhao, C. Goerg, and A. Timm-Giel, "Lte mobile network virtualization," *Mobile Networks and Applications*, vol. 16, no. 4, pp. 424–432, 2011.
- [14] 3rd Generation Partnership Project, "Etsi ts 136 413," Tech. Rep., 2013.
- [15] M. Yang, Y. Li, L. Zeng, D. Jin, and L. Su, "Karnaugh-map like online embedding algorithm of wireless virtualization," in *Wireless Personal Multimedia Communications (WPMC)*, 2012 15th International Symposium on, 2012, pp. 594–598.
- [16] 3rd Generation Partnership Project, "3GPP Technical Specification," Tech. Rep., 2014.
- [17] Cisco, "Cisco Visual Networking Index : Global Mobile Data Traffic Forecast Update 2014-2019," Tech. Rep., 2014.