

Challenges for 5G Transport Networks

Matteo Fiorani¹, Paolo Monti¹, Björn Skubic², Jonas Mårtensson³,
Luca Valcarenghi⁴, Piero Castoldi⁴ and Lena Wosinska¹

¹KTH Royal Institute of Technology, Electrum 229, SE-164 40 Kista, Sweden

²Ericsson AB, Färögatan 6, SE-164 80 Kista, Sweden

³Acreeo Swedish ICT AB, Isafjordsgatan 22, SE-164 40, Kista, Sweden

⁴ Scuola Superiore Sant'Anna, Via Moruzzi 1, 56124 Pisa, Italy

Abstract—5G mobile communications is seen as the enabler for the networked society where connectivity will be available anywhere and anytime to anyone and anything. The details of 5G are the subject to ongoing research and debate, mostly focused on understanding radio technologies that can enable the 5G vision. So far, less work has been dedicated to the challenges that 5G will pose to the transport network. This paper provides a first analysis of the key challenges to 5G transport in terms of capacity, flexibility and costs, for example. Different use cases are discussed as well as technology options and control plane concepts.

I. INTRODUCTION

The 5th generation of mobile networks (5G) is the next major phase of mobile telecommunications. There is yet no official specification of 5G, but rather a general consensus on the vision of what 5G should be able to achieve: a future with user- and machine-centric communications where access to information and sharing of data is available anywhere and anytime to anyone and anything, i.e., the so called Networked Society [1].

There are a number of initiatives at the national and international level that are currently working on 5G [2][3]. Among them the EU FP7 project METIS [4] defined a set of scenarios and use cases that future 5G networks are expected to support [5]. This list includes examples such as the dense urban information society, shopping mall, virtual reality office, vehicle communications, just to name a few of them. The support for these scenarios and test cases will require a new generation of mobile access beyond the mere upgrade of the existing ones, i.e., 3G and 4G. 5G will need to offer not only higher peak-rates per subscriber, but also handle a larger number of simultaneously connected devices, to provide better coverage and to offer highly reliable communications with low outage probability and latency. 5G must also achieve the above objectives at a similar cost and energy consumption as today's networks.

Some of the proposed radio technologies for 5G include new and flexible spectrum resources (e.g., sub-6.5GHz, mm-wave Radio Access Technologies RAT), multi-antenna technologies with massive multiple-input and multiple-output (MIMO) and beam-forming, as well as inter site coordination for higher spectral efficiency. Increased capacity and coverage will also be achieved through continued radio site densification (e.g., through small cells). 5G may also include different deployment architectures (e.g., integrated radio base station, centralized baseband, radio-over-copper solutions) depending on different deployment scenarios.

Most ongoing research efforts on 5G mobile networks concern radio technologies that will be part of 5G, whereas less is known on the implications on transport. However, with continued site densification and with larger numbers and variegated services to be provisioned, the role of the transport network, i.e., backhaul of radio base stations or fronthaul of remote radio units, will become more and more crucial.

Motivated by this last consideration this paper investigates the key challenges on transport with respect to future 5G mobile networks considering both small cell transport and metro aggregation. Based on the challenges for wireless access, as defined by the METIS project, the main transport challenges are defined. Possible solutions addressing these challenges are also presented and discussed. We find that the transport infrastructure for future 5G deployment needs to both accommodate high traffic volumes over specific geographical areas and provide transport resources in a flexible manner. This quality will be indispensable for operational purposes but also to limit as much as possible deployment costs without compromising the quality of experience of the various provisioned services.

II. 5G TRANSPORT CHALLENGES

In order to understand the 5G transport challenges one must understand how 5G may evolve the radio access segment. Among the various initiatives that are looking into 5G, the EU project METIS [5][6] defines 5G in terms of *scenarios* which the next generation wireless access networks will have to support. A total of five future scenarios have been defined, namely *amazingly fast* (users want to enjoy instantaneous network connectivity), *great service in a crowd*, *ubiquitous things communicating* (i.e., effective support to Internet of Things), *super real time and reliable connections*, and *best experience follows you*. Each of these scenarios introduces a challenge (see Fig. 1(a)). Three of these challenges (i.e., *very high data rate*, *very dense crowds of users* and *mobility*) are more traditional in the sense that they are related to continued enhancement of user experience and supporting increasing traffic volumes and mobility. Two emerging challenges, *very low latency* and *very low energy, cost and massive number of devices*, are associated with the application of wireless communications to new areas. Future applications may be associated with one or several of these scenarios imposing different challenges to the network. In METIS twelve specific test cases were defined and mapped onto the five scenarios (see Fig. 1(a)). The selected test cases essentially sample the space

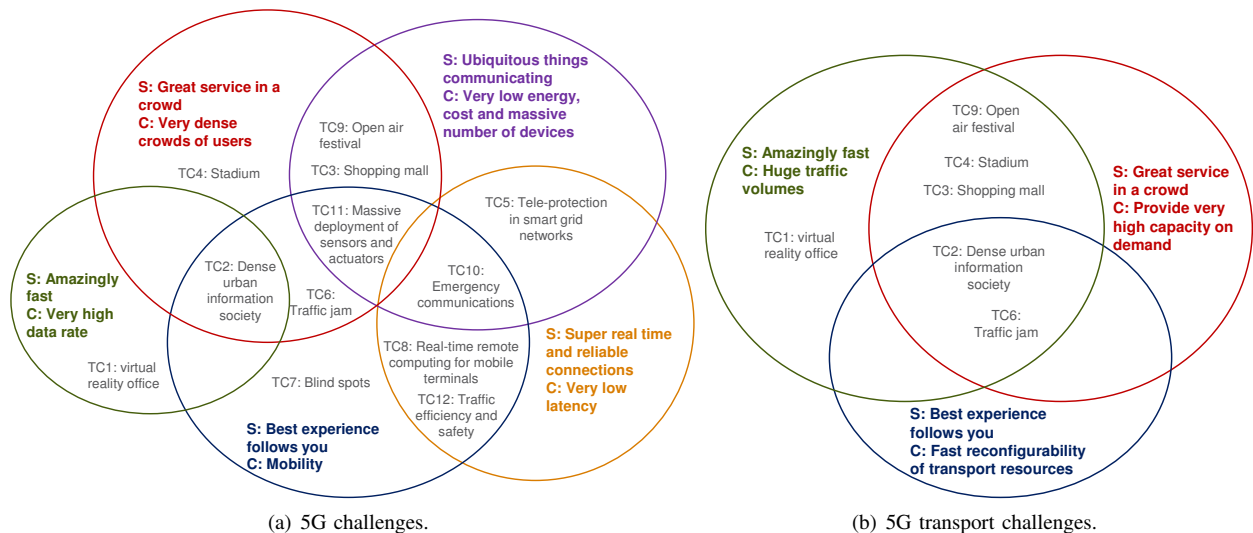


Fig. 1. 5G challenges (C), scenarios (S) and test cases (TC).

of future applications. Once technical enablers that fulfill the requirements for these test cases are defined, it is expected that other applications subject to the same fundamental challenges, will successfully be supported. As a consequence, defining technical enablers for the 5G test cases means also defining technical solutions to the 5G challenges.

While METIS is focused on wireless access, the challenges defined for 5G are expected to impact also the transport. Support for very high data rates will require both higher capacity radio access nodes as well as a densification of radio access sites. This, in turn, translates into a transport network that needs to support more sites and higher capacity per site, i.e. *huge traffic volumes*. The great service in a crowd will put requirements on the transport network to *provide very high capacity on-demand* to specific geographical locations. In addition, the best experience follows you scenario, suggests a challenge in terms of *fast reconfigurability of the transport resources*. On the contrary, the other 5G challenges are not expected to play as important role for shaping the transport, as for example the case of very low latency and very low energy, cost and massive number of devices. A properly dimensioned transport network based on modern wireless and/or optical technologies is already today able to provide extremely low latency, i.e., the end-to-end delay contribution of the transport network is usually almost negligible [7]. In addition, while a huge number of connected machines and devices will create a challenge for the wireless network, it will most probably not significantly impact the transport. This is due to the fact that the traffic generated by a large number of devices over a geographical area will already be aggregated in the transport. The three scenarios for the transport network described above are summarized in Fig. 1(b) along with their corresponding challenges and test cases. Note that Fig. 1(b) does not report all the original test cases but only those that pose challenges to the transport network. This information will be used later in the paper to identify the appropriate transport technologies.

III. CANDIDATE TECHNOLOGIES FOR 5G TRANSPORT

This section provides an overview of a number of transport options for 5G wireless networks. A 5G transport network can be divided in two different segments, i.e., *small cell transport* and *metro/aggregation* (Fig. 2). The small cell transport segment aggregates the traffic to/from the wireless small cells towards the metro/aggregation segment. Different solutions in terms of technology (e.g., optics, wireless) and topology (e.g., tree, ring, mesh) are possible [8] depending on the specific wireless access scenario. The metro/aggregation segment, on the other hand, connects different site types (i.e., macro and/or small cells) among themselves and to the core network, the latter via the service edge (service node for the interconnection among different network domains).

For the metro/aggregation segment one promising solution is represented by a dense-wavelength-division multiplexing (DWDM) - centric network [9][10]. In such a network, packet aggregation takes place at the edges of the network (e.g., at the small/macro cells sites and at the service edge), while at the center (i.e., between access and metro rings) switching is done completely in the optical domain thanks to active optical elements such as wavelength selective switches (WSSs) and reconfigurable optical add-drop multiplexers (ROADMs). It has already been demonstrated that DWDM-centric solutions have the potential to offer high capacity (in the order of tens to hundreds of Gbps) and lower energy consumption than their packet-centric counterparts (i.e., with packet aggregation at the center of the network) [11]. For this reason the DWDM-centric metro/aggregation concept may represent a good candidate for future 5G transport networks.

For the dedicated small cell transport segment it is not possible to define a single best candidate technology because of the variety of the small cell deployments which are test-case dependent. The main options can be categorized in copper-, fiber- and wireless-based technologies [12]. Wireless based solutions are attractive where cost of deploying wired transport infrastructure is prohibitively high. Copper-based options are able to offer rates in the order of a few Gbps over relatively




Fig. 2. 5G network segments.

TABLE I
OVERVIEW OF OPTICAL TECHNOLOGIES FOR SMALL CELLS TRANSPORT.

Technology	Compatible services	Capacity	CAPEX	OPEX	Maturity
Active Ethernet	Ethernet	1-100Gbps	Equipment (low/med), Deployment (low)	General (high), Energy (high)	Mature
PtP grey optical links	Ethernet, CPRI, RoF	1-100Gbps	Equipment (low/med), Deployment (high)	General (high), Energy (low/med)	Mature
TWDM-PON	Ethernet, (WDM overlay with CPRI/RoF over same fiber infrastructure)	$\leq 10\text{Gbps}$ (≤ 4 wavelengths in the feeder fiber)	Equipment (med), Deployment (low)	General (low), Energy (med)	Standardized
DWDM-PON	Ethernet, CPRI, RoF	1-10Gbps (≤ 380 wavelengths in the feeder fiber)	Equipment (high), Deployment (med)	General (low), Energy (med)	Products available for less cost sensitive segments. Low cost technologies in 3-5 years
UD-WDM-PON	Ethernet, CPRI, RoF	1-10Gbps (≤ 1000 wavelengths in the feeder fiber)	Equipment (high), Deployment (high)	General (low), Energy (med)	Concept/demo stage. Low cost technologies in 5+ years

short distances and therefore can be preferred in the areas where there is a large installed base of copper that can be reused. Optical transmission technologies are able to provide high data rates over long distances in an energy-efficient way. As a consequence, fiber-based solutions are seen as a good and long term candidate for 5G small cells transport networks [13]. Table I summarizes the main characteristics of some of the most promising ones. The table is organized in terms of: *transport services* that can be supported (e.g., Ethernet, common public radio interface (CPRI), and analog radio over fiber (RoF)); provided *capacity* per client site; *cost* (i.e., divided in CAPEX, which further is divided into equipment and deployment costs, and OPEX, where the energy cost impact is highlighted separately from the more general maintenance costs); and, finally, *technology maturity*.

For some test cases, like the open air festival, optical-based small cell transport may not be preferable, due to high deployment costs. In these cases, wireless-based solutions may represent a better alternative, mainly because they are usually easier, faster, and cheaper to deploy. Modern wireless transport technologies are able to provide very high capacities over short and medium distances, thanks to the introduction of new transmission paradigms, such as MIMO, and the opening of new spectrum ranges. Table II presents an overview of the most promising wireless small cell transport options. For each technology the table provides information about: the supported *transport services*; *capacity* and *reach*; sensitivity to *weather conditions*; *cost*; and *technology maturity*.

IV. ADDRESSING THE 5G TRANSPORT CHALLENGES

Section II discussed challenges for 5G transport, i.e., to accommodate huge traffic volumes, on-demand capacity, and the

ability to quickly reconfigure transport resources. This section presents a number of avenues to address these challenges.

One avenue is based on *overprovisioning* transport resources both in the small cell transport and the metro/aggregation transport segments, i.e., dimensioning for the worst case traffic scenario (i.e., highest possible traffic demand). The advantage of this approach is relatively low complexity since it does not require any dynamic transport control. With overprovisioning, the challenge of providing high capacity on-demand with (possibly) fast resource reconfiguration is satisfied thanks to the ubiquitous availability of ultra-high capacity transport. However, this solution potentially comes at high cost because of inefficient use of network resources.

A second avenue consists in a smarter network, employing advanced network functionalities (i.e., *dynamic resource sharing* and *network function virtualization (NFV)*) to quickly and efficiently respond to the transport needs. Dynamic resource sharing implies a sharing of transport network resources among the wireless base stations. The type of resources and how they can be shared depends on the specific 5G test case (e.g., traffic dynamicity), radio deployment solution/architecture and on the choice of transport technology. Examples of transport resources that can be dynamically shared are: optical fibers (e.g., using WDM), wavelength channels (e.g., using either time division multiplexing (TDM) or dynamic wavelength allocation), and spectrum (e.g., using time, frequency or code multiple access). NFV provides another dimension to flexibility through the ability to dynamically push network functions to different locations, such as closer to the users so that a portion of the traffic requests can be served locally. This can be used to offload network segments from traffic during congestion. Examples of network functions

TABLE II
OVERVIEW OF WIRELESS TECHNOLOGIES FOR SMALL CELLS TRANSPORT.

Technology	Compatible services	Capacity + Reach	Weather sensitivity	CAPEX	OPEX	Maturity
Microwave (6-42 GHz)	Ethernet	Max. 1Gbps; 5-10km	Low/med	Equipment (low), Deployment (low)	General (med), Energy (high)	Mature
E-band (71-76, 81-86 GHz)	Ethernet, CPRI	Max. 2.5Gbps; 1-2km	High	Equipment (high), Deployment (high)	General (med), Energy (high)	Products available for less cost sensitive segments. Low cost technologies in 3-5 years)
Free Space Optics (FSO)	Ethernet, CPRI, RoFSO	Max. 10Gbps; <1km	Very high	Equipment (high), Deployment (low)	General (high), Energy (low)	Products available for less cost sensitive segments. Low cost technologies in 3-5 years

that can be dynamically virtualized are baseband processing, evolved packet core, packet aggregation, and caching. One important aspect to consider for this avenue is the business model. If the transport and the wireless service provider are two separate entities, then the network functions that can be virtualized depend on the specific agreements between these two entities.

The use of dynamic resource sharing and NFV puts requirements on the control plane. A software-defined networking (SDN) [10] based control plane with programmable control of network resources and end-to-end orchestration could provide a framework for such a scenario. It could enable dynamic optimization of the use of transport network resources and provide a framework for interaction with other controllers (i.e., wireless network controller and IT controller) [14]. On the other hand, the design and implementation of such a complex control plane introduces several challenges. A main challenge is the definition of an orchestration entity able to keep track of the availability of different type of resources (e.g, wireless, transport and IT) and perform end-to-end optimization. A centralized orchestration entity leads also to scalability issues, which could be addressed through the adoption of a multi-layer control architecture and resources abstraction models. Different multi-layer SDN-based control architectures are possible depending on how the controllers of different segments are expected to interact. Fig. 3 shows an example of such a multi-layer controller. The dedicated small cells transport controller manages the resources in the dedicated small cells transport network and provides the overall transport controller with an abstract and simplified view of this network segment. The overall transport controller is in charge of managing all the transport resources and providing connectivity services to the other controllers (i.e., wireless and IT). The controllers interact via the orchestrator, which possesses an abstract and simplified view of all the resources and performs end-to-end provisioning and optimization.

V. CASE STUDIES

This section presents design of the transport network for two selected 5G test cases (i.e., *virtual reality office*, and *open air festival*) where dynamic resource sharing and NFV are considered.

A. Virtual Reality Office

A virtual reality (VR) office [6] is a top-modern (virtual) office space where the work involves the interaction with

high resolution 3D scenes. A team of 5 to 10 individuals simultaneously interact with the scene. Some of the team members are located within the same building, while others are working remotely from other office buildings (not necessarily in the same city/region). The majority of the traffic is generated by 3D video applications and the traffic volume density is as high as 0.1 Gbps/m^2 (assuming an average office size of 200 m^2 the traffic volume is 20 Gbps per VR office) [6]. In terms of wireless deployment the project METIS envisions an ultra dense small cell network (UDN) solution [6].

The overall impact of this test case on the metro network depends on how many VR offices an office building has. If we assume up to 50 VR offices then we can expect an aggregated backhaul traffic of up to 1 Tbps per building. In this scenario it is reasonable to assume that each building is directly connected to an access point of a fiber-based transport solution (e.g., a DWDM-centric metro network (see Fig. 4)), which is equipped with wavelength channels operating at high rate, e.g., 100 Gbps. Using a pure over-provisioning approach, the metro network should be dimensioned to reserve 10 wavelengths per building. As a consequence, in order to serve a metropolitan area with 100 office buildings a total of 1000 wavelength channels are required, leading to very high costs for the network operator. On the other hand, not all the offices in the area will run high capacity 3D video applications at the same time. Using a SDN control plane, such as the one proposed in Fig. 3 and an efficient dynamic wavelength allocation policy it is possible to substantially reduce the number of wavelengths required to cover the area. As an example, assuming that at busy hours only 60% of the VR offices will be generating high capacity 3D video traffic a total of 600 wavelength channels are theoretically sufficient to serve the area. It is then possible to save deployment costs, due to the lower number of transponders needed at the service edge, and energy consumption, by dynamically switching off the unused transponders inside the buildings. Additional benefits can also be achieved by placing IT resources servers at the access points or at internal nodes of the metro network (Fig. 4) and exploiting NFV. The servers can be controlled by the orchestrator, so that some network functions, such as caching of data files or processing of 3D video traffic, can be dynamically virtualized and pushed closer to the VR offices in order to save capacity inside the metro network. A possible policy for the orchestrator could be to monitor the traffic generated in each office building and to activate the local

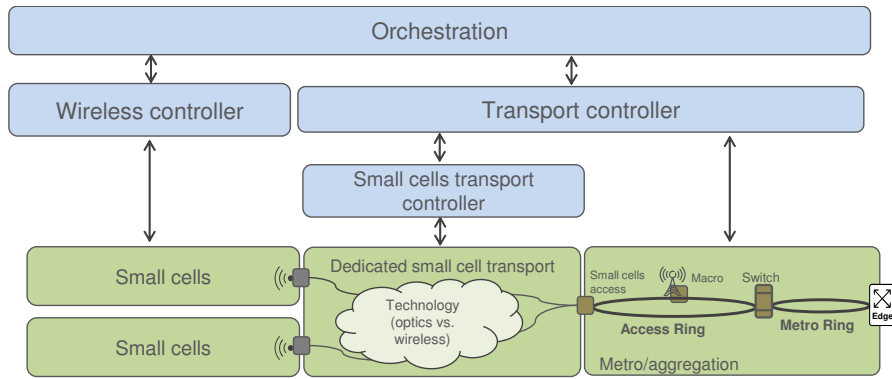


Fig. 3. Example of a SDN-based control plane architecture.

servers when the traffic exceeds a certain threshold. In this way the amount of traffic that needs to flow through the metro network can be reduced.

In terms of the dedicated small cell transport network, optical solutions are the most suitable given the high capacity needed between the wireless UDN and the metro aggregation point (Fig. 4). Assuming that VR offices will be deployed in modern buildings equipped with an advanced fiber infrastructure, the use of grey point-to-point (PtP) optical links might be an effective technology option due to low energy consumption and its relatively high maturity (Table I). In addition, a dedicated small cell transport controller can provide the operator with the possibility to reduce energy consumption and save costs by utilizing traffic variations and switching off and on the optical links accordingly.

B. Open Air Festival

The open air festival test case is defined as a small rural area visited by 100k visitors during a four-days-long multi-stage open air music festival [6]. During the event visitors may want, for example, to be able to share real-time or recorded high definition video clips, and to access the Internet at a high-speed, especially during the breaks between the performances. As a result the total aggregated traffic volume can be as high as 900 Gbps/km² over an area of around 1 km² (i.e., total traffic volume can be up to 900 Gbps) [6]. In terms of wireless deployment the project METIS envisions a temporary infrastructure consisting of 25 micro cells deployed in a grid with spacing 250 m [6] connected to a permanent macro base station, which is normally used to serve the considered rural area (Fig. 5).

It can be assumed that the permanent macro base station serving the rural area is directly connected to an access node of the DWDM-centric network (Fig. 5), while the traffic generated by the temporary micro cells is aggregated at the permanent macro and transmitted through the metro network. Aggregation networks are usually dimensioned to provide relatively small capacities in rural areas (due to the limited aggregated traffic requirements). Let's assume that each macro base station is normally equipped with a 10 Gbps wavelength channel. As a consequence, to satisfy the requirements of this test case, the capacity provided to the permanent macro covering the area of the open air festival should be temporarily

increased from 10 Gbps to 900 Gbps. A first solution, based on overprovisioning, would be to temporarily setup 9 dedicated optical links at 100 Gbps between the permanent macro base station and the service edge. This can be done by the transport operator when the temporary wireless infrastructure is set up, assuming also that fibers and wavelengths are already available because of the over-provisioning in the deployment phase. This approach is relatively simple and efficient, but also expensive (due to the deployment of the extra resources). On the other hand it has been already reported in many studies that macro base stations in rural areas are inactive most of the time [15]. When a macro base station is inactive, the corresponding wavelength in the metro network can be dynamically switched off or reallocated by a SDN-based control plane. In the case of the open air festival, all the wavelengths associated with inactive macro base stations can be moved to the macro located in proximity of the event. Assuming that the transmission capacity over these wavelengths can be dynamically upgraded from 10 Gbps to 100 Gbps, it is theoretically possible to provide enough metro capacity to serve the test case without need to overprovision the metro network in the deployment phase. This proposal is well in line with recent efforts from many operators to move from an "always on" (macro stations are always active) to an "always available" scenario (macro stations are activated on demand) [15].

Additional benefits can be obtained by exploiting NFV in this case as well. During the deployment of the temporary wireless infrastructure, the operator could place a certain number of servers at the macro base station site. These servers can be used to virtualize a number of network functions (e.g., virtual baseband processing, virtual evolved packet core) and for caching. In particular, caching is a very promising approach in this test case. In fact, the content requested by the users in an event such as an open air festival is expected to be relatively predictable. The operators can then prefetch some content in the caches in order to reduce the amount of traffic that needs to flow through the metro network. Based on simple calculations coming from the analysis of the cache ability of YouTube video traffic [16], the caching of YouTube video alone can lead to a reduction of around 20% of the total capacity generated by the open air festival (reducing the traffic volume from 900 to 720 Gbps).




Fig. 4. Transport network for the virtual reality office.




Fig. 5. Transport network for the open air festival.

A temporary dedicated small cells transport network is needed in this test case in order to connect the micro cells to the macro base station aggregation site (Fig. 5). Given the temporary nature of the solution, a wireless technology could be a better choice than an optical one. Among the wireless technologies, free space optics may be the best option, thanks to the high transmission capacities, but it is applicable only if the distance between the micro cells and the macro base station site is not too long. If the distance between the temporary wireless infrastructure and the macro is long, both microwave and E-band are promising technologies. Microwave has the advantage of cost and maturity, while E-band has the advantage of providing higher capacity (so that it could better support transport services such as CPRI).

VI. CONCLUSIONS

This paper provides an overview of the challenges to transport networks introduced by 5G. The three main transport challenges are identified: (i) huge aggregated traffic volumes, (ii) on-demand provisioning of very high capacity in specific geographical locations, and (iii) need for fast reconfigurability of the transport resources. Two approaches for designing and dimensioning a future 5G transport network have been considered. One is based on over-provisioning of transport resources while the second is based on dynamic resource sharing and network function virtualization (NFV) aided by a software defined network (SDN)-based controller. These two approaches have been compared in two specific 5G test cases, i.e., virtual reality (VR) office and open air festival. The analysis highlighted that the use of a common DWDM-centric access/metro network combined with a dedicated small cell transport can be an efficient choice for the future 5G transport. In addition, a SDN-based transport controller, able to efficiently perform dynamic resource sharing and NFV, helps in achieving high resource utilization and in reducing deployment costs.

VII. ACKNOWLEDGMENTS

The work described in this paper was carried out with the support of the Kista 5G Transport Lab (K5) project funded by VINNOVA and Ericsson, and of 5GrEEn ("Towards Green 5G Mobile Networks") project funded by EIT ICT Labs.

REFERENCES

- [1] Ericsson "Networked Society Essential" booklet, <http://www.ericsson.com>.
- [2] The 5G Infrastructure Public Private Partnership, <http://5g-ppp.eu/projects/>.
- [3] 3GPP The Mobile Broadband Standard, http://www.3gpp.org/news-events/3gpp-news/1614-sa_5g.
- [4] The METIS 2020 project: laying the foundation of 5G, <http://www.metis2020.com>.
- [5] METIS deliverable D1.1, "Scenarios, requirements and KPIs for 5G mobile and wireless system", April 2013.
- [6] METIS deliverable D6.1, "Simulation guidelines", October 2013.
- [7] B. Skubic, P. Öhlen, A. Rostami, Z. Ghebretensae, N. Lindqvist, K. Larraqi, "Optical Transport and Challenges in the Networked Society", to appear in Photonic Network Communications.
- [8] P. Monti, S. Tombaz, L. Wosinska, J. Zander, "Mobile Backhaul in Heterogeneous Network Deployments: Technology Options and Power Consumption", in Proc. of ICTON, July 2012.
- [9] Shuqiang Zhang, Ming Xia, S. Dahlfort, "Fiber routing, wavelength assignment and multiplexing for DWDM-centric converged metro/aggregation networks," in Proc. of ECOC, Sept. 2013.
- [10] P. Öhlen, B. Skubic, Z. Ghebretensae, W. John, M. Shirazipour, "Software-defined networking in a multi-purpose DWDM-centric metro/aggregation network," in Proc. of GLOBECOM, December 2013.
- [11] B. Skubic, I. Pappa, "Energy consumption analysis of converged networks: Node consolidation vs metro simplification," in Proc. of OFC, March 2013.
- [12] S. Tombaz, P. Monti, F. Farias, M. Fiorani, L. Wosinska, J. Zander, "Is backhaul becoming a bottleneck for green wireless access networks?," in Proc. of ICC, June 2014.
- [13] F. Farias, P. Monti, A. Vstberg, M. Nilson, J. Costa, L. Wosinska, "Green Backhauling for Heterogeneous Mobile Access Networks: What Are the Challenges?," in Proc. of ICICS, December 2013.
- [14] The open networking foundation (ONF), <https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/optical-transport-use-cases.pdf>.
- [15] P. Frenger, P. Moberg, J. Malmudin, Y. Jading, I. Godor, "Reducing Energy Consumption in LTE with Cell DTX," in Proc. of VTC, May 2011.
- [16] L. Braun, A. Klein, G. Carle, H. Reiser, J. Eisl, "Analyzing Caching Benefits for YouTube Traffic in Edge Networks - A Measurement-based Evaluation," in Proc. of NOMS, April 2012.