CHARTING THE FUTURE OF INNOVATION VOLUME 92 | #8 • 2015

ERICSSON TECHNOLOGY Review

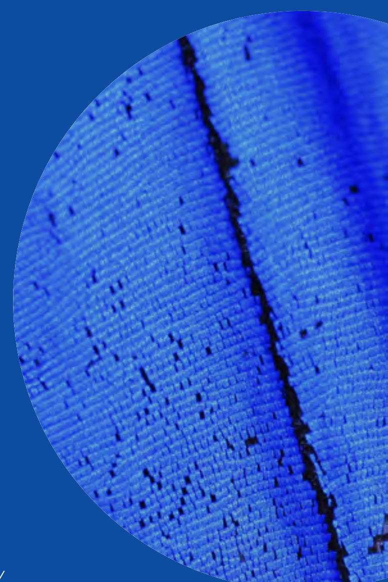

FLEXIBILITY IN 5G TRANSPORT NETWORKS: THE KEY TO MEETING THE DEMAND FOR CONNECTIVITY

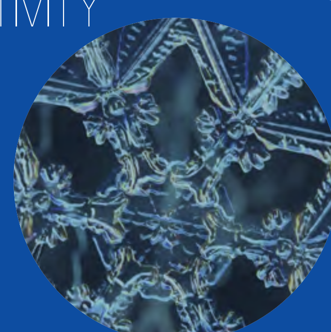

ERICSSON

FLEXIBILITY IN 5G transport networks

THE KEY TO MEETING THE DEMAND FOR CONNECTIVITY

The more people have been able to achieve while on the move, the more dependent society has become on mobile broadband networks. As applications like self-driving vehicles and remotely operated machinery evolve, become more innovative, and more widespread, the level of performance that 5G networks need to deliver will inevitably rise. Keeping pace with ever-increasing demand calls for greater flexibility in all parts of the network, which in turn requires tight integration between 5G radio, transport networks, and cloud infrastructures.

PETER ÖHLÉN
BJÖRN SKUBIC
AHMAD ROSTAMI
KIM LARAQUI
FABIO CAVALIERE
BALÁZS VARGA
NEIVA FONSECA LINDQVIST

ADVANCES IN TECHNOLOGY and a shift in human behavior are influencing how 5G networks are shaping up. With 3G, things got faster, data volumes surpassed voice, new services were developed, and people started using mobile broadband. With 4G, mobile broadband soared. Today's networks provide advanced support for data. Building on this success, 5G aims to provide unlimited access to information and the ability to share data anywhere, anytime by anyone and anything. So, as we move deeper into the Networked Society, the connections that link things and people will become almost exclusively wireless.

Services like mobile broadband and media distribution will continue to evolve in line with our growing global dependence on connectivity. Networks will experience huge increases in traffic and will need to service an ever-expanding number of connected devices – both massive MTC (IOT) and mission-critical MTC. The latter sets stringent requirements for performance characteristics like reliability and latency.

The digital and mobile transformations currently sweeping through industries worldwide are giving rise to innovative cross-sector applications that are demanding in terms of network resources. And so, 5G networks will not only need to meet a wide range of requirements derived from user demand and device development; they will also need to support advanced services – including those yet to be developed.

Limitless innovation in application development, device evolution, and network technology are shifting from a model that is operator steered to one that is user driven. Flexibility and operational scalability are key enablers for rapid innovation, short time to market for deployment of services, and speedy adaptation to the changing requirements of modern industry.

How will future networks evolve?

To ensure that networks will be able to cope with the varied landscape of future services, a variety of forums like NGMN, ITU-R, and 5G PPP are working on the definition of performance targets for 5G systems [1].

In comparison with 2015 levels, the performance projections that will have most impact on transport networks are:

- **1000x mobile data volume per geographical area, reaching target levels of the order of Tbps per sq km**
- **1000x the number of connected devices, reaching a density of over a million terminals per sq km**
- **5x improvement in end-to-end latency, reaching to as low as 5ms – as is required by the tactile internet.**

However, the maximum levels of performance will not all apply at the same time for every application or service. Instead, 5G systems will be built to meet a range of performance targets, so that different services with widely varying demands can be deployed on a single infrastructure.

Getting networks to provide such different types of connectivity, however, requires flexibility in system architecture.

Aside from meeting the stringent requirements for capacity, synchronization, timing, delay, and jitter, transport networks will also need to meet highly flexible flow and connectivity demands between sites – and in some cases even for individual user terminals [2].

Emerging 5G radio capabilities and the convergence of radio access and wireless backhaul have triggered an uptake of fixed wireless technologies as a complement to fixed broadband [3]. With hybrid access 5G networks will be able to provide the increased capacity needed to handle peak traffic for residential users. As such, 5G radio will increasingly complement and overlap with traditional fixed-broadband accesses.

Terms and abbreviations

5G PPP–5G Infrastructure Public Private Partnership | API–application programming interface | BB–baseband | CPRI–Common Public Radio Interface | CWDM–coarse wavelength division multiplexing | DWDM–dense wavelength division multiplexing | EPC–Evolved Packet Core | FTTH–fiber-to-the-home | MIMO–multiple-input, multiple-output | MPLS–multi-protocol label switching | MTC–machine-type communication | NFV–Network Functions Virtualization | NGMN–Next Generation Mobile Networks | NG-PON2–next-generation passive optical network | P router–provider router | PE router–provider edge router | PGW–PDN gateway | ROADM–reconfigurable optical add/drop multiplexer | SDN–software-defined networking | SLA–Service Level Agreement | UE–user equipment

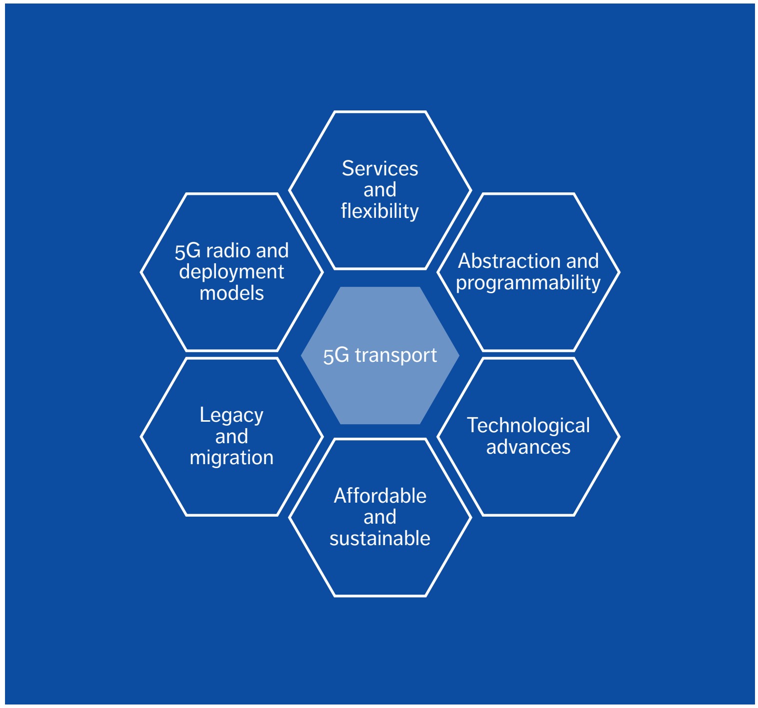

Figure 1
Landscape for 5G transport

The 5G transport network

As 5G radio-access technologies develop, transport networks will need to adapt to a new and challenging landscape, as illustrated in *Figure 1*.

Services

The expectations for 5G networks are high – providing support for a massive range of services. Industry transformation, digitalization, the global dependence on mobile broadband, MTC, the IOT, and the rise of innovative industrial applications all require new services, which has a considerable impact on the transport network. For example, a new radio-access model that supports highly scalable video distribution or massive MTC data uploading might require additional transport facilities – such as a scalable way to provide multicasting.

5G radio

How the 5G radio is deployed determines the level of flexibility needed in the transport network. Capacity, multi-site and multi-access connectivity, reliability, interference, inter-site coordination, and bandwidth requirements in the radio environment place tough demands on transport networks.

In 5G, traditional macro networks might be densified, and complemented through the addition of small cells. Higher capacity in the radio will be provided through advances in radio technology, like multi-user mimo and beamforming, as well as the availability of new and wider spectrum bands [4]. Consequently, the capacity of the 5G radio environment will reach very high levels, requiring transport networks to adapt. Not only will transport serve a large number of radio sites, but each site will support massive traffic volumes, which might be highly bursty due to the peak rate available in 5G.

For example, a UE that is connected to a number of sites simultaneously, may also be connected to several different access technologies. The device may be connected to a macro over LTE, and to a small cell using a new 5G radio-access technology. Multi-site and multi-RAT connectivity provides greater flexibility in terms of how UEs connect to the network and how E2E services are set up across radio and transport. For example, allowing for efficient load balancing of UEs among base stations not only improves user experience, it also improves connection performance.

The impact of interference may favor deployment models where coordination can be handled more effectively. In small-cell deployments, UEs are often within reach of a number of base stations, which increases the level of interference, and at times requires radio coordination capabilities for mitigation. However, the method used for handling interference depends on how transport connectivity is deployed. In a centralized baseband deployment, tight coordination features, such as joint processing, can be implemented. In traditional Ethernet and IP-based backhaul, tight coordination requires low-latency lateral connections between participating base stations.

Centralized baseband processing tends to result in lower operational costs, which makes this approach interesting. However, it typically comes at the cost of high CPRI bandwidths in the transport network. The high bandwidth, together with stringent delay and jitter requirements, makes dedicated optical connectivity a preferred solution for fronthaul.

In 5G networks, the bandwidth requirements for fronthaul could be very high. The demand will be created by, for example, antennas for MU-MIMO and beamforming – which could use in the order of 100 antenna elements at each location. In combination with dense deployments and wider frequency bands (in the 100MHz range) traditional CPRI capacity requirements can quickly reach levels of several Tbps. A new split of RAN functionality is under investigation to satisfy requirements for cost-effective deployments and radio performance, while keeping capacity requirements on transport within a manageable range.

But some primary networking principles remain valid, such as timing and synchronization. Defining new packet-based fronthaul and midhaul interfaces requires the underlying network to include protocols and functions for time-sensitive transport services. Related standardization efforts are currently underway [5].

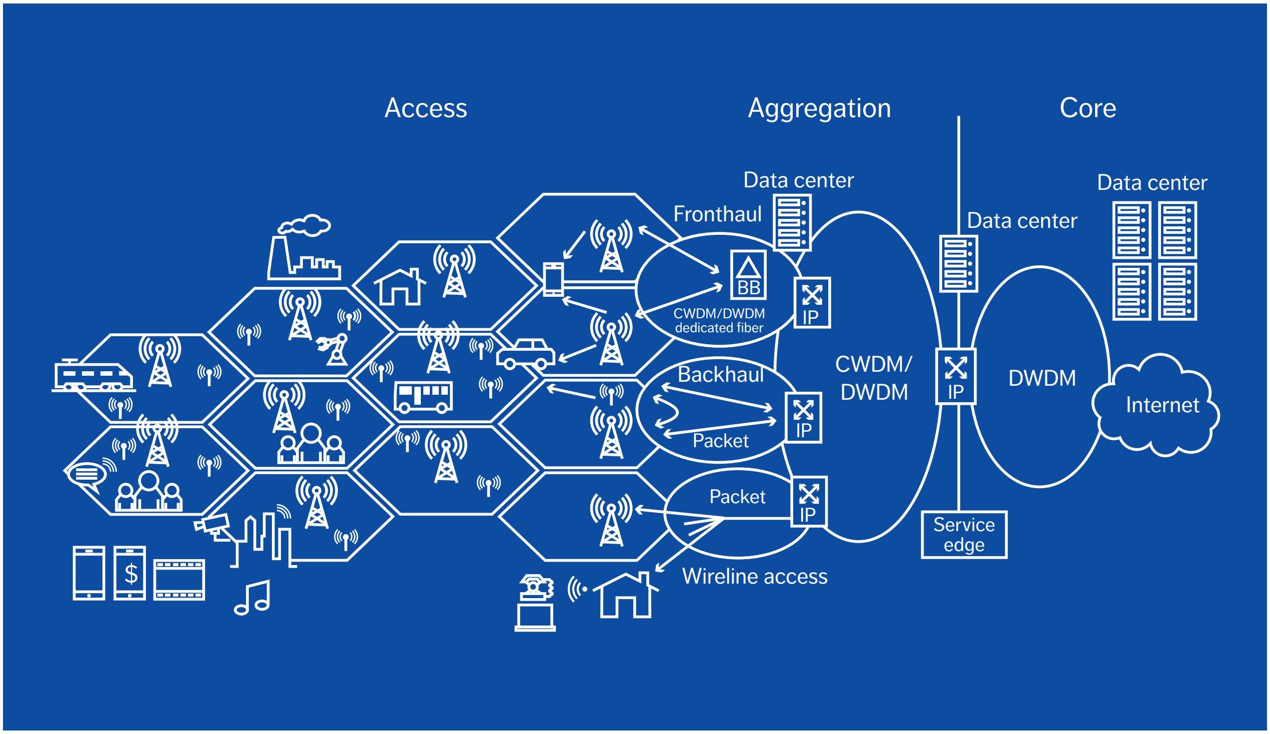

Figure 2
Main technology options to connect RAN and transport infrastructure

Abstraction and programmability
Abstracting network resources and functionality, as well as managing services on-the-fly through programmatic APIs are the pillars of SDN, and the source of its promise to reduce network complexity, and increase flexibility.

With a new split in the RAN, some functions can be deployed on general-purpose hardware, while others, those closer to the air interface with strict real-time characteristics, should continue to be deployed on specialized hardware. Most of the functions of the EPC will be deployed as software – following the concept of Network Functions Virtualization (NFV). Deploying network functions in this way makes it possible to build end-to-end network slices that are customized for specific services and applications. Each layer of the network slice, including the transport layer, will be designed to meet a specific set of performance characteristics.

The significance of network slices is best illustrated by comparing applications with different requirements. A network of sensors, for example, requires the capability to capture data from a vast amount of devices. In this instance, the need for capacity and mobility is not significant. Media distribution, on the other hand, is challenged by large capacity requirements (which can be eased through distributed caching), whereas the network characteristics for remote-control applications based on real-time video are high bandwidth and low latency.

From a 5G-transport perspective, there is a need to provide efficient methods for network sharing, so that applications like these – each with their individual requirements, including mechanisms to satisfy traffic isolation and SLA fulfillment – can be supported for several clients. In addition, distributed network functions need to be connected over links that fulfill set performance levels for bandwidth, delay, and availability.

Transport networks will need to exhibit a high degree of flexibility to support new services. To this end, key features are abstraction and programmability in all aspects of networking – not just connectivity but also storage and processing.

Legacy, migration, and new technologies

The main technologies that contribute to performance enhancement and the network segment – access, aggregation, or core – they apply to are outlined in ***Figure 2***. 5G transport will be a mix of legacy and new technologies. Long-term, network evolution plans tend to include fiber-to-the-endpoint. In practice, however, providing small-cell connectivity requires that local conditions be taken into consideration, which results in the need for several technologies – such as copper, wireless links, self-backhauling, and free-space opto – to be included in the connectivity solution. Re-use of existing fixed access infrastructure [6] and systems will be important, and new technologies and systems may in turn provide more efficient use of available infrastructure. For example, additional capacity can be provided by extending the use of CWDM and DWDM closer to the access segment of the network. At the same time, interworking with IP is essential to provide end-to-end control, and to ensure that the fiber infrastructure is used efficiently.

Existing infrastructure, together with operator preferences, determines the necessary evolution steps, and how the migration process from legacy to desired architecture should proceed.

> “PROGRAMMABILITY IN 5G TRANSPORT NETWORKS WILL IMPROVE FLEXIBILITY”

The design of 5G transport networks will need to continue to be affordable and sustainable, keeping the cost per bit transported contained. Handling legacy in a smart way, and integrating sustainable advances in technology into packet and optical networks will help to keep a lid on costs.

Programmable control and management

Flexibility through programmability is a significant characteristic that will enable 5G transport networks to support short time to market for new services and efficient scaling.

Programmability gears up networks, so they can take on innovations rapidly, and adapt to continuously changing network requirements. A couple of capabilities need to be determined to enable programmability for transport networks:

- **the required degree of flexibility or ability to reconfigure**
- **the layer or layers that need to be programmable.**

Determining these capabilities is a trade-off between need and gain; in other words, how does the benefit of programmability compare with the cost of the technology needed to provide it? A significant factor for transport providers in weighing up need against gain is how to address packet-optical integration. This is because extending programmability to the optical layer not only provides greater flexibility and ease of provisioning to allocate transport bandwidth; it also simplifies the process of offloading the packet layer through optical/router bypass, as well as providing improved cross-layer resilience mechanisms [7].

The telecom industry has long set itself two principal targets for transport networks: efficient resource utilization, and dynamic service provisioning and scaling. While these goals still stand, they need to be revised continually to match the changing needs of client layers. These needs include the short reaction times demanded by modern applications, and the fact that different clients will need to interface with the network at different layers. Add connection capabilities like bandwidth and latency into the mix, and the need for network programmability becomes more evident.

So just how does increased network programmability help the telecom industry meet the targets it has set for itself, given the need for different performance characteristics for different applications?

Efficient resource utilization

Transport programmability enables network operators to exploit traffic dynamicity to optimize the utilization of resources across different segments of the network.

A programmable transport network facilitates the division of transport resources into multiple (isolated) slices. These slices can be allocated to different clients – enterprises or service providers – enabling efficient sharing of resources.

Dynamic service provisioning and scaling

Being able to provision resources on the fly is particularly crucial for dynamic service chaining, which involves interconnecting distributed, virtualized network functions and ultimately facilitating dynamic service creation. In particular, establishing connection services across several networking domains has long been a challenge – here enhanced programmability can make such procedures more efficient. In most cases, flow control in the transport domain should be carried out on aggregated traffic to avoid detailed steering for individual users when it is not needed.

A programmable transport network enables the capacity allocated to a service to be scaled up or down, when and where it is needed across the network – in other words, providing elastic services.

Centralized or distributed control

Control plays an essential role in programmability. Network control can be centralized or distributed, and networks are operated differently depending on the approach used.

Centralized control – the concept used in SDN – enables shorter service development cycles and speedier rollout of new control functionality (implementation occurs once in the central stack). For networks built with a distributed control plane, changes must be made in multiple – already deployed – control stacks (especially in multi-provider networks).

The topic of SDN is being discussed in the telecom industry as a promising toolset to facilitate network programmability. In SDN architecture, the main intelligence of network control is decoupled from data plane elements and placed into a logically centralized remote controller: the SDN controller (SDNC). As such, the SDNC provides a programmatic API, which exposes abstracted networking infrastructure capabilities to higher layer control applications and services, enabling them to dynamically program network resources.

The role of the API in SDN goes beyond traditional network control. It allows applications to be deployed on top of the control infrastructure, which enables resources to be automatically optimized across heterogeneous network domains, and new end-to-end services to be instantiated easily. The control/management system needs to provide methods for controlling resources and for exposing infrastructure capabilities – using the right abstraction with the level of detail suitable for higher layer applications.

To highlight this point, in our research we chose to exemplify the case of resource and service orchestration across multiple network domains with heterogeneous types of resources. The resulting hierarchical SDN-based control architecture, which orchestrates across three domains – transport, radio access networks (RANs), and cloud – is shown in *Figure 3*. A management function [8], which can be partly overlapping, is included but not discussed in detail in this article.

SDN flavors

The impact of upgrading the control plane of a legacy transport network to SDN depends on a number

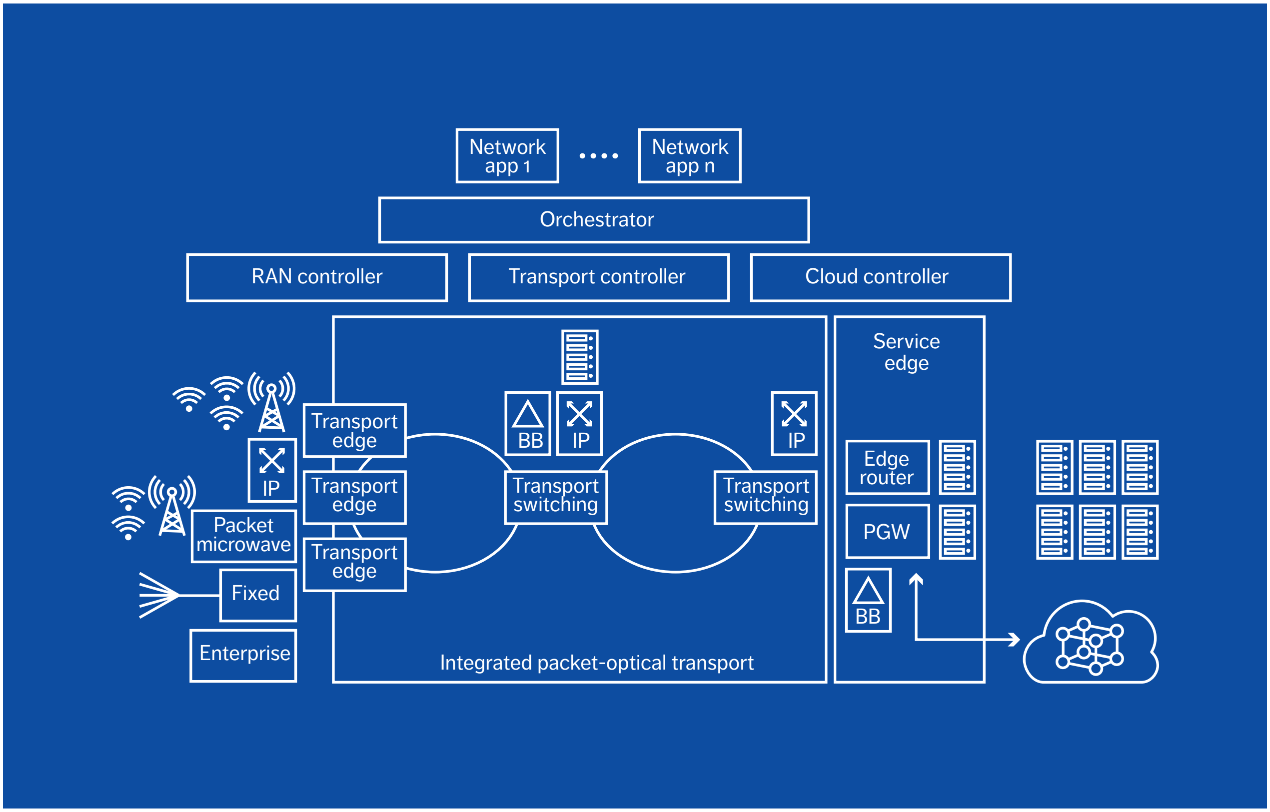

Figure 3
Hierarchical SDN control architecture for multi-domain orchestration

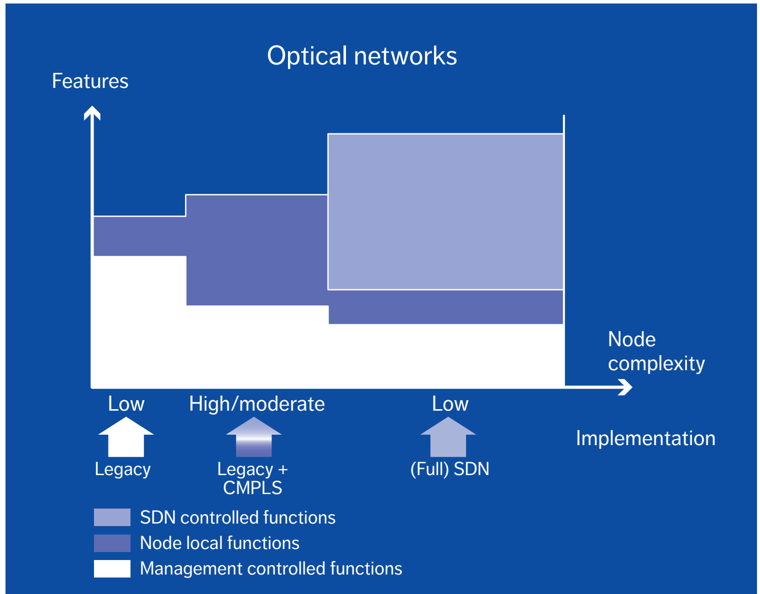

Figure 4a
Centralizing control functionality in the optical domain

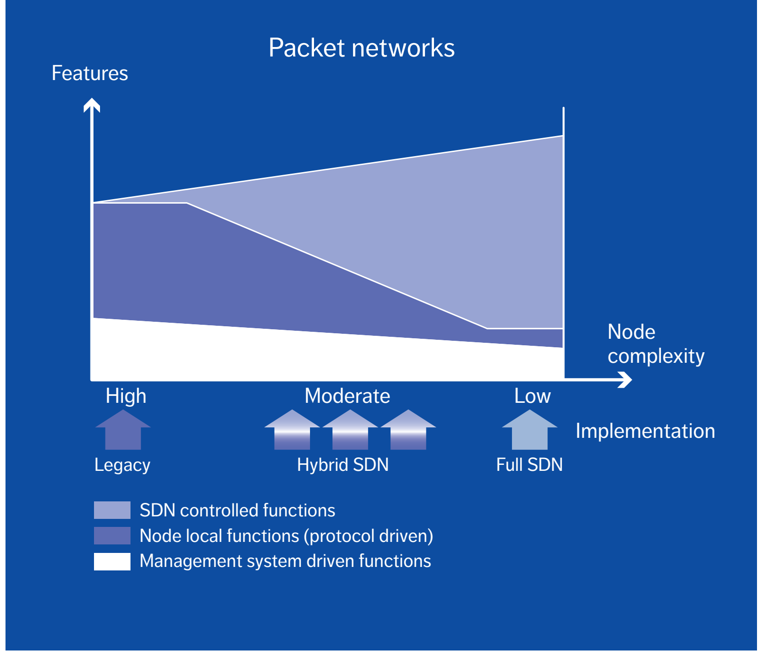

Figure 4b
Centralizing control functionality in the packet domain

of aspects, but primarily on the degree to which forwarding and control functions are integrated in legacy transport networks.

In legacy optical transport networks, most control functions are already separated from the data plane nodes. However, in packet-switched transport nodes, the two planes are tightly coupled. As such, introducing a fully centralized SDN control plane (full SDN) is more straightforward for optical transport than packet networks.

To integrate legacy optical transport networks, as illustrated in *Figure 4a*, the SDNC needs to be developed along with suitable interfaces, but it does not necessarily require disruptive changes to the optical nodes.

When applied to packet networks, the disruption created by SDN is significant. A more natural approach for the packet domain would be to centralize selected elements of the control functions. The resulting hybrid-SDN alternatives are illustrated in *Figure 4b*. Ideally, control over service-related functions should be centralized with the SDNC, while transport-related functions should be implemented locally on the node.

The decision of where to place a control-plane function or feature is operator specific, depending on many factors, like the available feature set, and operational preferences.

In packet-based transport networks, the concept of separating transport- and service-related functions is well established, where a clear logical differentiation is made between service unaware transport nodes and service nodes – such as P and PE routers in MPLS networks. Only service nodes hold service states and require implementation of service-related functions. Such separation is a future-proof concept and one that should remain intact. Any improvements in this area should focus on transport service functions, as they cause most of the challenges in building and operating networks and make the introduction of new services lengthy and costly – especially in multi-vendor environments.

Flexible transport plane

Several factors contribute to network dynamicity – the ability of a network to adapt rapidly to changing demand. The introduction of 5G radio technologies and the launch of new services are the two main factors pushing the need for networks to be more dynamic, and consequently the need for a more flexible transport plane. There are, however, many other factors contributing to network dynamicity:

- **resource dynamicity: on-the-fly addition and removal of connectivity, compute, and storage resources**
- **traffic dynamicity: responsiveness to fluctuating traffic patterns that result from user movement/migration, or variations in user activity [9]**
- **service dynamicity: responsiveness to service usage patterns with widely varying resource requirements**
- **failures and service windows: ability to reroute traffic and minimize impact of downtime**
- **weather conditions: managing the effect of rain or fog on performance for microwave or free-space opto networks.**

Expert opinion differs on how access and transport architectures will evolve to meet future mobile requirements, and in particular how they will provide support for small cells. Legacy networks typically consist of separate branches for fixed (residential/business services) and mobile access. Continued densification of mobile networks is likely to result in the use of several different small-cell transport technologies – each one adapted for specific network conditions.

In particular, the adoption of wireless backhaul/fronthaul technologies, such as NLOS, for provisioning connectivity to new small-cell sites, will become more prevalent. At the same time, the fixed-access infrastructure and its widespread availability will continue to be valuable for providing small-cell connectivity, pushing the need for fixed and mobile network convergence.

How fixed/mobile convergence might evolve depends on existing fixed access infrastructure. Many operators have been reluctant to invest in deep fiber technologies, like FTTH, due to the costs associated with deployment, and have instead turned to alternatives like copper-based drop links using DSL or CAT5/6. This type of architecture or active optical network relies on the presence of a large number of distributed active nodes.

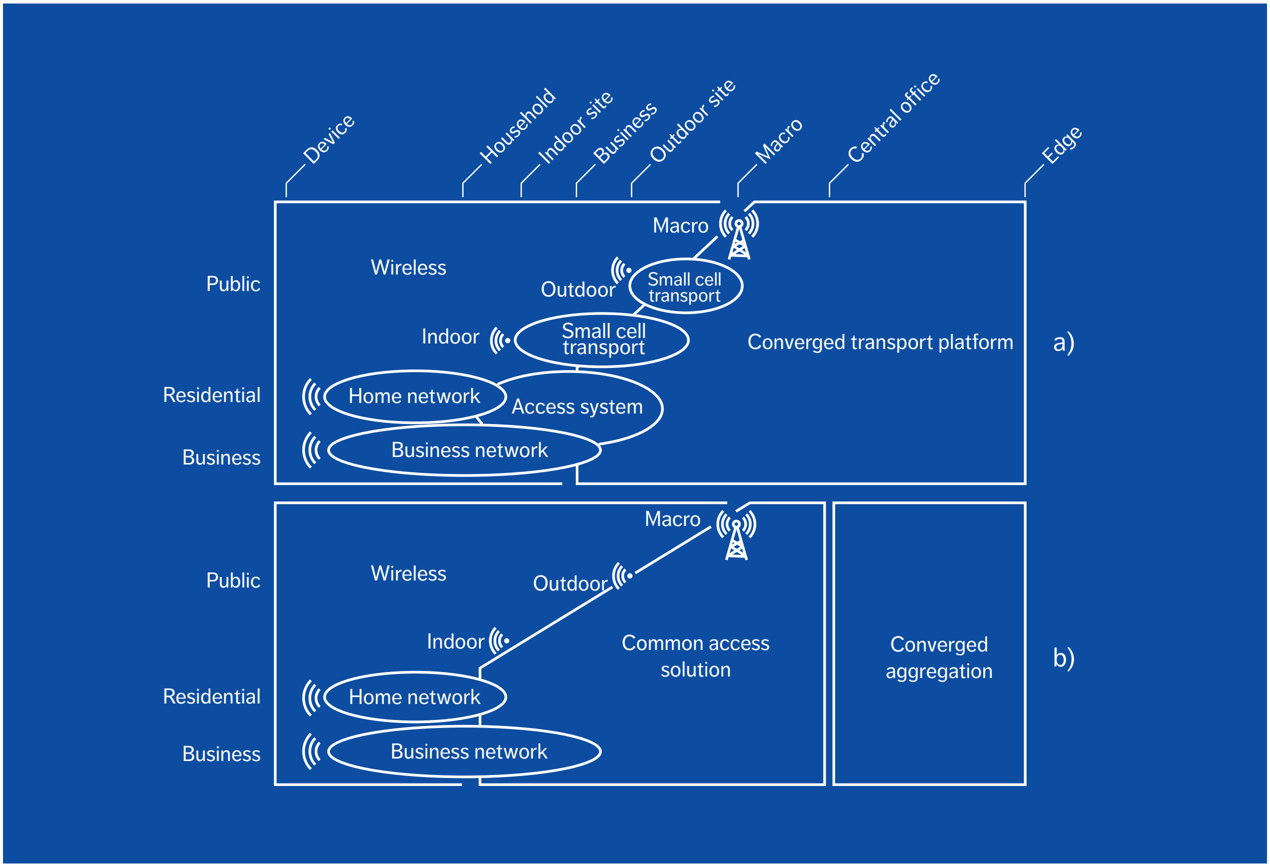

Figure 5
Different architectural avenues with scenarios based on:
(a) a converged transport platform
(b) a common access solution

Figure 5 illustrates two evolution scenarios for fixed/mobile convergence. Different options are available for providing converged access infrastructure for traditional residential and business access services, as well as IP-based backhaul and CPRI-based fronthaul [6]. In the bottom part (b) of the illustration, the connectivity needs of the RAN are served through a common access solution. Here, the challenge is to define a system that can simultaneously meet the cost points of residential access and the performance requirements for different RAN deployments. As illustrated in the top half (a) of Figure 5, evolving networks in this manner paves the way for a converged transport platform (possibly including small-cell fronthaul) under common control but with a somewhat diverse data plane, subject to requirements for different segments. In reality, many more scenarios and combinations exist than are illustrated here, and the choice of architectural infrastructure model, RAN deployment model, and technology are all closely interlinked. In turn, the decisions made for network architecture and technologies determine the degree of flexibility that transport can offer, whether it is at the packet or wavelength level.

In a converged scenario, one possible solution is to deploy nodes capable of providing common switching of packet and fronthaul – which today are separate domains that use different transport protocols (Ethernet and CPRI) with their own specific requirements. To multiplex Ethernet and CPRI, switching at wavelength and packet layers can be combined: the challenge is to meet latency and jitter requirements for time-sensitive applications. Deterministic delay switching using client agnostic frames is an alternative to packet for scenarios where existing DWDM/OTN metro infrastructure is used simultaneously for backhaul and fronthaul applications.

For optical interfaces, fiber rich deployments can benefit from recent developments of grey 100G and 400G optical interconnection interfaces. This is an active field of research, and many solutions, standardized and proprietary, are being explored. In some, modulation formats are natively designed to dynamically adjust the bandwidth according to real-time networking needs. Integrated photonic technologies are changing the optical communications industry outlook, promising a dramatic reduction of hardware cost, power consumption, and footprint, but also enabling flexibility at lower cost. Multi-channel transceivers and tunable lasers are among the first applications targeted by integrated photonics, as demonstrated by standardization activity on NG-PON2 and G.metro. Along the same technological trend, Ericsson has defined a new type of device – an integrated reconfigurable optical add-drop multiplexer tailored to the specific needs of radio access. Such a device would provide flexibility in the wavelength layer and would be an order of magnitude simpler and cheaper than conventional ROADMs.

Conclusions

The first 5G network trials are already ongoing on a small scale, and commercial systems are expected in 2020. Comparing 5G with previous generations shows that it is not just a new radio-access technology – so much more is expected of it. 5G is shaping up to provide cost-effective and sustainable wireless connectivity to billions of things, people, enterprises, applications, and places around the world. To make the most of this business opportunity and deliver connectivity to billions of devices, the architecture of 5G systems – and transport in particular – needs to be built for flexibility through programmability.

Delivering the required level of flexibility, needs tighter integration between 5G radio, transport networks, and cloud infrastructure. This must be carried out with a backdrop of small-cell deployment, convergence of access and backhaul, and migration of legacy equipment and technologies – while containing costs.

When it comes to programmability, the expectations placed on SDN technology to deliver are enormous. SDN also brings service velocity with it, as well as a means to integrate transport, radio, and cloud domains. However, adopting a hybrid-SDN alternative might be best to mitigate the disruption SDN causes when applied to packet networks. This approach enables control to be centralized for service functions and distributed for transport – gaining a degree of flexibility without the disruption.

To meet capacity demands, increasing the use of DWDM closer to access will be feasible when flexible optics become more cost-efficient. In short, the major challenges for 5G transport are programmability, flexibility, and finding the right balance of packet and optical technologies to provide the capacity demanded by the Networked Society.

References

1. **5G PPP, March 2015, 5G vision: the next generation of communication networks and services, available at:** *https://5g-ppp.eu/wp-content/uploads/2015/02/5G-Vision-Brochure-v1.pdf#page=8*
2. **NGMN Alliance, February 2015, NGNM 5G White Paper, available at:** *https://www.ngmn.org/uploads/media/NGMN_5G_White_Paper_V1_0.pdf*
3. **Ericsson Review, November 2014, Wireless backhaul in future heterogeneous networks, available at:** *http://www.ericsson.com/news/141114-wireless-backhaul-in-future-heterogeneous-networks_244099435_c*
4. **Ericsson Review, June 2014, 5G radio access, available at:** *http://www.ericsson.com/news/140618-5g-radio-access_244099437_c*
5. **IEEE, 802.1CM - Time-Sensitive Networking for Fronthaul**
6. **COMBO FP7 project: Convergence of fixed and Mobile Broadband access/aggregation networks, more information can be found on:** *http://www.ict-combo.eu*
7. **Ericsson Review, May 2014, IP-optical convergence: a complete solution, available at:** *http://www.ericsson.com/news/140528-er-ip-optical-convergence_244099437_c*
8. **Ericsson Review, Nov. 2014, Architecture evolution for automation and network programmability, available at:** *http://www.ericsson.com/news/141128-er-architecture-evolution_244099435_c*
9. **Ericsson, May 2015, Research Blog, Transport, Radio and Cloud Orchestration with SDN, available at:** *http://www.ericsson.com/research-blog/5g/transport-radio-and-cloud-orchestration-with-sdn/*

THE AUTHORS

Peter Öhlén

◆ is a principal researcher within IP and transport, managing research efforts across network and cloud domains. His current research focuses on

transport network solutions for 5G, and integrating heterogeneous technology domains – transport network, radio and cloud. He joined Ericsson in 2005, and has worked in a variety of technology areas such as 3G radio, HSPA fixed-wireless-access, fiber access, IP and optical networks. He holds a Ph.D. (2000) in photonics from the Royal Institute of Technology (KTH) in Stockholm, Sweden.

Björn Skubic

◆ is a senior researcher in IP and transport, and is currently managing

activities in 5G transport. He joined Ericsson in 2008, and has worked in areas such as optical transport, energy efficiency and fixed access. He holds a Ph.D. in physics from Uppsala University, Sweden.

Ahmad Rostami

◆ is a senior researcher at Ericsson Research, where he manages activities in the area of programmable 5G transport networks and radio-transport-cloud orchestration. Before joining Ericsson in 2014, he worked

at the Technical University of Berlin (TUB) as a senior researcher and lecturer. At the university, his areas of interest covered SDN as well as the design and performance evaluation of broadband communication networks. In 2010, he received a Ph.D. from the TUB Faculty of Electrical Engineering and Computer Science for his work on control of optical burst switched networks. Rostami holds an M.Sc. in electronic engineering from Tehran Polytechnic, Iran.

Neiva Fonseca Lindqvist

◆ is a senior researcher in RAN transport solutions, with a particular interest in heterogeneous networks and the evolution to 5G. She joined Ericsson in 2011, working with fixed and mobile backhaul

network architectures. Previously, she worked in research in the area of signal processing for broadband communication over copper access networks and held a postdoctoral position at the EIT– LTH Faculty of Engineering, Lund University, Sweden. Fonseca Lindqvist holds a Ph.D. in electrical engineering and telecommunication from the Federal University of Para (UFPA), Belém, Brazil.

Kim Laraqui

◆ is a principal researcher in mobile backhaul and fronthaul solutions for heterogeneous networks. He joined Ericsson in 2008 as a regional senior customer solution manager. Prior to this, he was a senior consultant on network solutions, design, deployment, and operations

for mobile and fixed operators worldwide. He holds an M.Sc. in computer science and engineering from KTH Royal Institute of Technology, Stockholm, Sweden.

Fabio Cavaliere

◆ joined Ericsson Research in 2005. He is an expert in photonic systems and technologies, focusing on WDM metro solutions for aggregation and backhauling networks and ultra-high-speed

optical transmission. He is the author of several publications, and is responsible for various patents and standardization contributions in the area of optical communications systems. He holds a D.Eng. in telecommunications engineering from the University of Pisa, Italy.

Balázs Varga

◆ joined Ericsson in 2010. He is an expert in multiservice networks at Ericsson Research. His focus is on packet evolution studies to integrate IP, Ethernet, and MPLS technologies for converged

mobile and fixed network architectures. Prior to joining Ericsson, Varga worked for Magyar Telekom on the enhancement of its broadband services portfolio and introduction of new broadband technologies. He has many years of experience in fixed and mobile telecommunication and also represents Ericsson in standardization. He holds a Ph.D. in telecommunication from the Budapest University of Technology and Economics, Hungary.

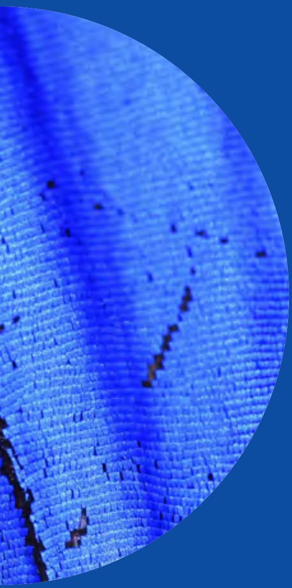

ISSN 0014-0171
284 23-3260 | Uen

Ericsson
SE-164 83 Stockholm, Sweden
Phone: + 46 10 719 0000