Towards Smart Public Interconnected Networks and Services – Approaching the Stumbling Blocks

Håkon Lønsethagen [1], Stanislav Lange [2], Thomas Zinner [1], Harald Øverby [1], Luis M. Contreras [1], Nicola Ciulli [1], and Emmanuel Dotaro [1]

[1]Affiliation not available
[2]NTNU

October 30, 2023

Abstract

The visions of 5G and beyond (B5G) imply unprecedented expectations toward general digital high-performing connectivity services in both public and private networks. Connectivity services that offer performance guarantees along multiple QoS dimensions are partially available today, but are confined to (virtual) private network services. However, open and equal access to public and Internet-scale Specialized Connectivity Services (SCS) delivered on-demand with interconnections across networks and support for mixed traffic modes that go beyond traditional best effort does not exist. In this paper, we argue that this is a huge industrial and societal problem that needs a solution. However, this problem is highly complex and multi-faceted, and there are many reasons why we are more or less locked into the status quo. This paper identifies the stumbling blocks and argues that there is a need for a holistic set of solution elements to take us across these hurdles. A roadmap approach is suggested alongside related research topics. Models and simulations are provided showing how a mixed traffic mode approach provides anticipated benefits. Arguments are given why the context brought by B5G will put us into conditions for change, allowing public SCS eventually at a global scale.

Towards Smart Public Interconnected Networks and Services – Approaching the Stumbling Blocks

Håkon Lønsethagen, Stanislav Lange, Thomas Zinner, Harald Øverby, Luis M. Contreras, Nicola Ciulli, Emmanuel Dotaro

***Abstract*—The visions of 5G and beyond (B5G) imply unprecedented expectations toward general digital high-performing connectivity services in both public and private networks. Connectivity services that offer performance guarantees along multiple QoS dimensions are partially available today, but are confined to (virtual) private network services. However, open and equal access to public and Internet-scale Specialized Connectivity Services (SCS) delivered on-demand with interconnections across networks and support for mixed traffic modes that go beyond traditional best effort does not exist. In this paper, we argue that this is a huge industrial and societal problem that needs a solution. However, this problem is highly complex and multi-faceted, and there are many reasons why we are more or less locked into the status quo. This paper identifies the stumbling blocks and argues that there is a need for a holistic set of solution elements to take us across these hurdles. A roadmap approach is suggested alongside related research topics. Models and simulations are provided showing how a mixed traffic mode approach provides anticipated benefits. Arguments are given why the context brought by B5G will put us into conditions for change, allowing public SCS eventually at a global scale.**

***Index Terms*—B5G, QoS, Specialized Connectivity Services, Public Networks.**

I. Introduction

The current ecosystem of Public Interconnected Networks and Services (PINS) is undergoing disruptive changes. These changes include the architectural transformations related to (B)5G, new requirements from a variety of stakeholders, and the paradigm shifts toward edge computing and network softwarization. New applications with stringent resilience and sub-millisecond latency [1] requirements are emerging. Hence, the degree of heterogeneity of services and applications that co-exist and compete for resources on a shared physical infrastructure is increasing. Moreover, the inadequacy of today's best-effort Internet to cope with such heterogeneity gives rise to solutions that rely on private networking and illustrates an increasing demand for beyond best-effort modes for public networks [2] that are more sustainable than and can complement traditional overprovisioning [3], [4].

As a result, network operators are facing challenges in multiple directions: Resources are required dynamically to cope with increasingly varying demands, systems become less predictable due to the heterogeneity of services and applications, and control plane complexity will grow with multi-dimensional Service Level Agreements (SLAs). Furthermore, service delivery expectations are growing from the vertical sectors, the online application providers (OAPs), and the consumer side towards a set of smart and specialized connectivity services on-demand (SCS) that are universally and equally provided. SCS must enable on-demand end-to-end connectivity with performance offerings w.r.t. multiple QoS dimensions. However, if the Internet remains as-is and SCS are totally isolated from the Internet, e.g., in the form of private dedicated networks, there will be a non-optimal use of resources and likely inefficient and costly business processes. Despite numerous efforts to address these aspects individually, no holistic solution has emerged so far and unsustainable short-term mitigations such as overprovisioning prevail.

Moreover, the (B)5G visions are changing the business and regulatory context of future telecommunications. We now have expectations from verticals, the public sector, and governments that (B)5G shall be smart and contribute to the digitalization in a green and efficient way. This calls for a new and holistic approach that can bootstrap SCS to become a fundamental connectivity layer that goes beyond today's best-effort Internet, and that can evolve and support smart PINS at a global scale. The grand challenge is to ensure that SCS are delivered across public interconnected networks between any end-point on the Internet, or where one or more of the end-points are located in a private network domain. An SCS instance will belong to a given logical network, and security-enabled *Logical Network as a Service* is anticipated to grow in volume and variety and dominate PINS. Note as well that [5] is complementary to our work, and early work in this direction has been conducted in the context of the NetWorld 2020 community [6].

Such a holistic approach needs to tackle technical, business-related, and regulatory challenges that are particularly hard to address in the traditional way since network operators tend to be locked-in to the status quo on both a technical and non-technical level. In this work, we provide a taxonomy of these challenges (Section II), and identify key solution elements and service concepts (Section III) that we argue are needed in order to overcome them. The overall desired outcomes or value propositions are as follows. i) Increased resource utilization, ii) increased energy efficiency, iii) predictable Quality of Experience (QoE) and customer utility, and perhaps even more importantly iv) unleashing a new generation of innovation potentials, in particular by SMEs that can not offer specialized application services (SAS) by building their own global backbone networks. In Section IV, we provide illustrative model- and simulation-based numerical examples of the expected benefits by allowing differentiated connectivity sharing common resources. We discuss the appropriateness, timeliness, and feasibility of the proposed approach, and conclude the article in Section V.

II. Stumbling Blocks, Uncertainties, and Challenges for Network Service Providers (NSPs)

NSPs are facing numerous and great uncertainties on both technical and non-technical issues, including business and user demand, privacy, and regulatory aspects.

A. Technological Challenges

The main technological challenge relates to the asymmetries w.r.t. the information flow. These include the applications', networks', and users' lack of ability to express their respective intents, offerings, and expected Quality of Experience (QoE) levels. It also includes the lack of corresponding interfaces to exchange these types of information within and between administrative domains. Moreover, establishing a reliable mapping from the QoS delivered by the network and the resulting QoE, which varies per specific application type, elasticity, and configuration, poses a significant challenge. Hence, networks are neither aware of the exact demands nor about the actually delivered QoE, resulting in barriers towards automated control loops that are required to efficiently and quickly adapt to dynamically changing network conditions. In particular, interaction between applications and the network should take place in both directions in order to enable application-aware networking as well as network-aware applications. Widespread use of encryption combined with plans for extending encryption towards transport headers further complicate efforts regarding such interactions [7].

Despite ongoing efforts for mapping KPIs to QoS characteristics as well as supplying isolated virtual networks (VNs) or slices that can deliver the corresponding QoS profiles [8], there is a lack of connectivity services that can be set up on demand with properties that are tailored to the characteristics of a given application and the needs of a specific user.

B. Business and Demand Uncertainties

On the non-technical side, business- and privacy-related as well as regulatory challenges arise. First, the predominant *overprovisioning cycle* of incremental capacity upgrades becomes increasingly unprofitable and unsustainable [3] due to ever stricter performance demands and therefore lower resource utilization. While the demand uncertainty is severe, a larger stumbling block is the fact that the business model for SCS is completely lacking. Solving this problem is a daunting multi-actor coordination challenge. Who wants to take the risk and effort to lead and facilitate the tasks that are needed, while at the same time facing the challenges and uncertainties around net neutrality (as discussed below)?

Furthermore, the introduction of SCS offerings has the potential to disrupt existing business models such as telco voice and enterprise VPNs, making operators hesitant towards entering into such new business models.

C. Privacy Challenges

In terms of privacy, both end-users and operators might have a limited willingness to share information, especially as long as it is not clear which level of detail this information shall be at and what kind of cost and / or performance gains can be expected in exchange.

The type of shared information can be organized with respect to three dimensions. First, device-related information that can range from a full specification to a detailed or just high level set of features. Similarly, application- or service-related information can include the specific application, the generic application class / type, or even just a high-level description of its needs and QoS sensitivities. Finally, operators involved in establishing end-to-end connectivity services could either exchange detailed information regarding their network topology and devices, aggregated information on available protocols and technologies, or just high-level information on possible interconnection options with other domains.

D. Regulatory Uncertainties

Differentiated treatment of network traffic needs to comply with regulations and in particular Network Neutrality (NN) principles. The current NN regulations put down in e.g., EU regulation 2015/2120 limit the network provider's options to offer differentiated classes of Internet access to end-users. Specifically, the introduction of differentiated traffic classes should not affect the quality of the (baseline) Internet Access Service (IAS).

The NN regulations are technology neutral, meaning that they also apply for 5G networks and beyond. However, it may be that (B)5G technology offers functionality that is legally limited by current NN regulations, thus prohibiting the full exploitation of the business potentials of (B)5G in an optimal manner.

We believe NN regulations should evolve to allow SCS in an open and neutral way with equal access for all, recognizing the separation between SCS and Specialized Application Services (SAS). Otherwise, a great amount of risk and uncertainty will prevail for NSPs, ISPs, and backbone Internet service providers.

III. Service Concepts, Main Solution Elements, and Recommendations

Considering the above, in this section we discuss requirements and main characteristics of service concepts and key principles around SCS to enable an evolution towards smart PINS. Furthermore, we identify key elements of the networking ecosystem that need attention to address the above challenges related to the long-term success of the service concepts, as well as a roadmap outlining a step-wise evolution in the short, medium, and long term. Our approach is inspired by "removing complexity and aiming towards simplicity".

A. Service Concepts and Overall Principles

Our discussion of service concepts revolves around two aspects that relate to the treatment of network traffic, namely traffic modes as well as traffic aggregates and connectivity handling. In the following, we define and discuss these notions in detail.

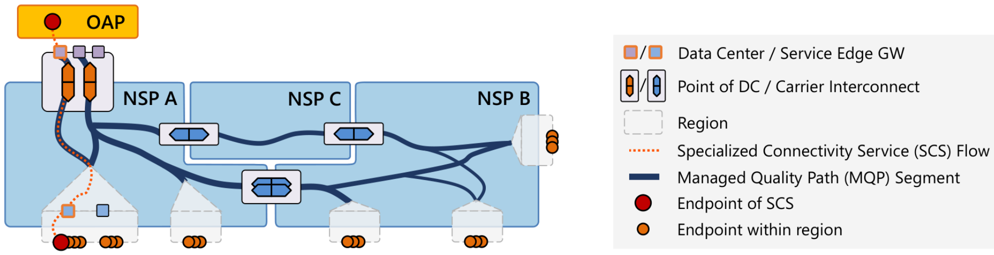

Fig. 1. Multi-domain scenario: managed quality path infrastructure with exemplary specialized connectivity service.

1) Traffic Modes: Due to the variety of applications in terms of their QoS demands as well as degrees of sensitivity and elasticity, we argue that diverse traffic modes which reflect this heterogeneity are required for efficient traffic handling. Depending on the tolerable amount of complexity, particularly control plane complexity, and with respect to the available technical solutions, relative differentiation between flows as well as absolute differentiation with strict performance guarantees can be performed. This way, QoS resources can be adjusted to reduce queueing delays for delay-sensitive traffic while identifying more delay-tolerant portions of the traffic that might be re-routed via longer paths. Evolving from the "best-effort" traffic mode of today's Internet as starting point, we exemplarily discuss a total of four traffic modes. These include three best-effort modes which differ relatively to each other and allow for more nuanced differentiation while retaining the benefits of best-effort handling, as well as an assured quality mode that provides strict performance guarantees.

From today's perspective, the current best-effort mode could be labeled "Basic Quality" (BQ). Relative to the BQ mode, we suggest an "Improved Quality" (IQ) mode. The IQ mode improves the quality or performance relative to the BQ mode by mechanisms such as Weighted Fair Queueing (WFQ). This may result in improvements potentially along multiple dimensions of the connectivity, in principle any combination of improved throughput, (queueing) delay, jitter, or packet-loss performance. Moreover, we foresee a "Background" (BG) traffic mode which provides connectivity with more relaxed quality properties than the BQ mode. This mode enables the NSP to provide valuable connectivity offerings while allowing a higher level of utilization of network resources. Examples for applications using the BG, BQ, and IQ mode could be the automatic download of an OS update, a user-initiated file download, and an on-demand video stream, which have increasingly strict QoS requirements while not necessarily being mission-critical and requiring strict performance guarantees.

Since throughput is less of an issue in backbone networks, the primary difference among the three best-effort modes is along the delay performance, where IQ offers lower end-to-end latency than the BQ mode, and the BG mode supports more relaxed latency requirements. All together, the BG, BQ, and IQ traffic modes can be considered as *multi-level best-effort*.

The fourth suggested traffic mode is referred to as Assured Quality (AQ) mode and offers strict performance guarantees. This mode is used if the client is in need of significantly higher or more stable network performance than available by the IQ mode. While the throughput available via the IQ offering is anticipated to be similar or somewhat higher than the BQ mode, the AQ mode is anticipated to enable significantly higher performance than what is available via the IQ mode. This will require mechanisms that are more complex than those of the IQ mode. That is, to enable the AQ mode, the QoS, resource, and admission control mechanisms must be realized at a finer granularity than those for the IQ mode.

2) Traffic Aggregates and Scalable Connectivity Handling: We observe that provisioning connectivity resources for each individual traffic flow in an on-demand and end-to-end fashion is not feasible in terms of scalability, complexity, and timeliness. Hence, we suggest introducing multiple granularity levels of traffic aggregates that differ w.r.t. their size, lifetime, and mode of instantiation. In this work, we discuss a two-level example. At the coarse-grained level, we propose *Managed Quality Paths (MQPs)* that are high-capacity, long-lived, and pre-established paths between major interconnection points. At the fine-grained level, we envision *SCS Sessions* that are expected to be highly dynamic, on-demand, and between endpoints. In the context of these sessions, only paths connecting the endpoints to suitable interconnection points need to be provisioned whereas the remainder can be carried by a suitable, already provisioned, and well-dimensioned MQP. These pre-established paths also help reducing the size of the solution space and therefore allow for faster handling of connectivity requests. Depending on the specific requirements of an SCS, different traffic modes might be employed at both the MQP and SCS session level. Figure 1 presents the main concepts alongside key infrastructure elements that are elaborated further in the following. We also show an exemplary SCS (flow) between the two highlighted endpoints to illustrate the core ideas.

Peering and transit services in today's Internet connect very large, coarse-grained regions and have only rudimentary SLAs. Evolving from these services, we propose the MQP as a *Point-of-Interconnect-to-Region* (PoI2R) interconnection service, where the notion of "region" can be in a spatial / geographical or technological sense, e.g., a range of IP prefixes. The core idea of the MQP service is to enable dynamic traffic engineering, intelligent management, and configuration of coarse traffic aggregates and their services, and it may also support remote peering (remote adjacencies).

Since SCS sessions use MQP for PoI2R connectivity, only paths from end-users and application servers to endpoints of

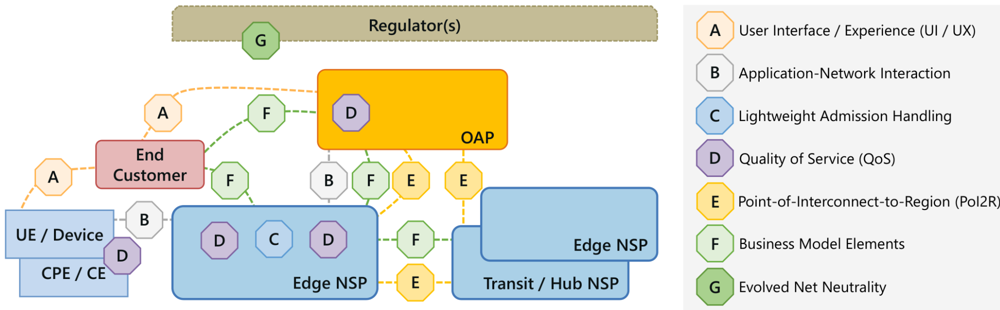

Fig. 2. Overview of solution elements.

the corresponding MQP need to be provisioned. We refer to these endpoints as Data Center Gateway (DC GW) and Service Edge Gateway (SEG), respectively. To achieve QoS handling and charging support for SCS between end-user devices and OAP end-points in data centers, we expect the necessity for signaling and business relationships between OAPs, NSPs, and end-user devices. With the proposed two-level approach, we expect that setting or merely checking policies at the gateways is sufficient to cover SCS needs. Hence, signaling and QoS handling can be substantially simplified in comparison to mechanisms such as IntServ that require setting policies on each network element along the entire end-to-end path.

B. Main Solution Elements and Challenges

In this subsection, we identify key aspects of the networking ecosystem that require attention in order to pave the way towards smart PINS and enable the discussed service concepts. We also highlight related challenges and some research questions (labeled *RQ*). Following the principles introduced above, we cover seven key topic areas that all need to be addressed in the long term and in a holistic, coordinated, and interdependent way. A compact overview of these areas and corresponding solution elements is provided in Figure 2. While we discuss all seven areas from A to G, we put particular emphasis on three of them that we consider critical already in the bootstrapping phase. These include Application-Network Interaction (ANI, B), Lightweight and Class-based Admission Handling (LAH, C), and Business Model Elements (BME, F).

UI / UX: The customer shall be in overall control of the service level selection. This can be achieved through explicit control via User Interface (UI) and User Experience (UX) dialogues or implicitly based on relevant characteristics of the environment and end devices in use. For this, we expect the need for a common approach to service level expectations and indicators (*RQ 1*).

ANI: The goal of ANI is to allow expressing requests for SCS and corresponding NSP offerings. To this end, the NSP could provide SCS templates expressing possible QoS value ranges - e.g., target and lower bound - in terms of supported throughput, latency, and packet loss. A solution should anticipate that elasticity of applications can still be an important feature, and support (re-)negotiation in case the desired quality level can not be delivered (*RQ 2*). Ideally, the SCS templates and the APIs for ANI should be standardized to ensure scalability, portability, and efficiency for the application developer. In order to guide and assist the application's creation of the SCS request, the ANI should also enable applications to query the network for available SCS profiles that can be delivered on demand or at / for a specific time. This way, the application can make requests with a high likelihood of success. Similar application-initiated reservation strategies have been proposed for SDN-based networks in the context of participatory networking [9].

LAH: According to the SCS session handling introduced above, we suggest a light-weight admission handling approach still sufficient to achieve the needed traffic handling targets. A fundamental and key objective is to ensure that the volume of admitted SCS sessions will not make the BQ mode and general quality level suffer beyond specific committed network performance levels. The concept of class-based admission handling is proposed and can be based on trust, logical network IDs, and per-class SCS treatment. Still, the NSP can monitor application traffic rate and behavior per OAP and can accordingly perform class-based and per-OAP policing by various means in a scalable way. Numerous QoS and admission control mechanisms are available at different levels of granularity in the access parts of 4G and 5G, and can be aligned with the overall end-to-end approach.

QoS: QoS mechanisms are available at multiple locations within the UE / device, edge NSP, and OAP domains. These mechanisms allow (re-)aligning the allocation of available resources with users' requirements.

PoI2R: Several variants of the PoI2R / MQP interconnection service are needed to cover the NSP-OAP and NSP-NSP segments that form SCS. In the context of NSP-NSP segments, an NSP is referred to as transit NSP if it does not contain any endpoint of the SCS.

BME: Along with the ANI and PoI2R service variants and enablers, BME and charging principles should be defined. Here, we underline the need to support also hybrid money flows from the customer to both the OAP and eNSP, and to optionally support Initiating Party Network Pays (IPNP) for two-way connectivity across NSPs. Note that at the PoI2R (MQP) level, the BME would be based only on traffic aggregates, and should support settlement-free and transit-only BMs. Moreover, we anticipate that the pricing models towards the end-customers in particular will evolve to ensure a correspondence between the price paid and the resources used, hence incentivizing responsible use of resources and "Green ICT".

Evolved Net Neutrality (eNN): Finally, we anticipate the need for evolving the Net Neutrality Regulation (eNN) to ensure that open and equal access to SCS can be supported at the scale of the global Internet. This will be critical in order to realize the vision of smart PINS and to achieve the desired *multi-level best-effort* traffic modes. A key question in this context revolves around expressing and measuring the general performance level of the BQ mode so as to ensure protection of the basic Internet access service (*RQ 3*).

C. Design Directions, Bootstrapping, and Roadmap

It will not be possible to create one perfect complete solution addressing all the above elements in just one go. A roadmap analysis is important, aiming first at addressing the low-hanging fruits, demonstrating that cost can be managed well, and that the effects are achievable. Addressing the solution elements in a trial & pilot context for bootstrapping is a critical short term step. Considering the European context in the following, this could be regarded as part of the Horizon Europe Smart Networks and Services (SNS) Joint Undertaking Partnership [10] and the anticipated Pan-European interconnected facilities for large-scale trials. Such a platform could allow experimentation and research to address the range of solution elements identified above, and exemplified in Figure 3.

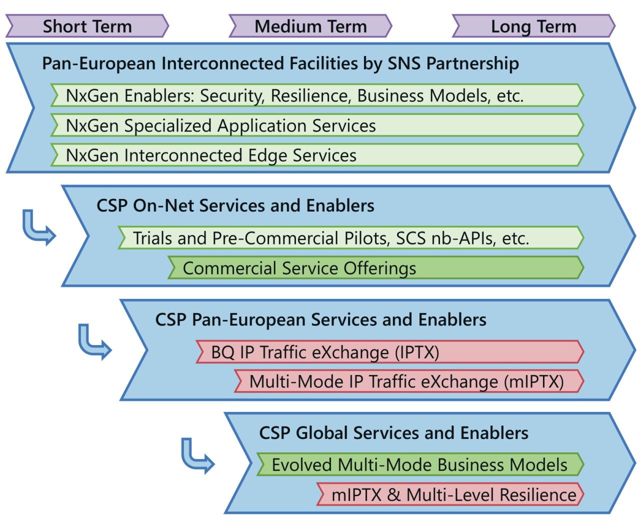

Fig. 3. Roadmap towards smart PINS. Topics are marked light green, darker green, and red to indicate a predominantly technology-, business-, and IPTX-oriented focus, respectively.

This stage can evolve towards bootstrapping SCS trials and pre-commercial pilots by the Communication Service Providers (CSPs). Pan-European trials and pilots will be important to provide evidence of the capabilities and measurement techniques that can enable and constitute the evolved net-neutrality solution elements, and will be key in the dialogue with regulators.

When both the solution elements as well as the business and regulatory context are sufficiently mature, the CSP market can evolve into the stage of "CSP Pan-European Services and Enablers". We suggest that the initial phase of this stage should focus on enabling QoS in the access segments, as these are the predominant bottlenecks. Many of the European CSP core networks and Internet IP Traffic eXchange (IPTX) peering links are already well-dimensioned and will be sufficient for this early stage. Hence, by just using the currently available BQ (aka. Best-effort) IPTX, a new generation of SLAs are achievable.

However, after a successful introduction of SCS offerings and service adoption in the market, a new generation of multi-mode IPTX is anticipated. This can be enabled by recognizing just a few DiffServ Code Points at the PoIs. In the long term, this approach may evolve into an universal solution enabling smart PINS.

IV. Evidence of Potential

In order to illustrate the potential of the proposed SCS ecosystem, we present results from several simulation-driven experiments, using the methodology from [11] and amending it with what-if analyses towards emerging applications, thereby focusing on principal effects.

The heterogeneity of applications is illustrated in Figure 4 where the QoS in terms of available bandwidth and experienced delay is plotted against the resulting QoE for a total of four applications. In addition to a TCP-based adaptive video streaming application and a UDP-based Voice over IP (VoIP) application, we perform what-if analyses involving applications with significantly stricter QoS requirements to mimic emerging and more demanding applications such as ones involving haptic feedback [12]. To this end, we use the QoS characteristics of the VoIP application, increase its bandwidth requirements and usage by a factor of 10, and increase its delay sensitivity by factors of 10 ("Delay Sensitive x10") and 20 ("Delay Sensitive x20"), respectively.

Several observations can be made. First, the relationship is clearly application-specific, both in terms of the sensitivity towards QoS changes as well as in terms of the absolute requirements. Second, applications exhibit different degrees of elasticity: while the video streaming application smoothly transitions between video quality levels, the VoIP application under consideration is inelastic and abruptly degrades from a good or acceptable QoE level to the worst level as soon as its bandwidth requirement is not met. Third, diminishing returns occur, i.e., past a certain application-specific point, lowering an application's delay or increasing its bandwidth slows and eventually stops improving its QoE. Additional QoE impact factors such as jitter, packet loss, and subjective user-specific properties are deliberately omitted in this work in order to keep a clear focus on the general ideas of the proposed concepts.

In the following, we highlight how the proposed multi-level, differentiated treatment of applications alongside optimized resource allocation can improve the situation for end-users, application providers, and network operators over a baseline best-effort approach. To this end, we individually pair up the adaptive video streaming application with each of the three delay-sensitive applications and assess the resulting QoE as

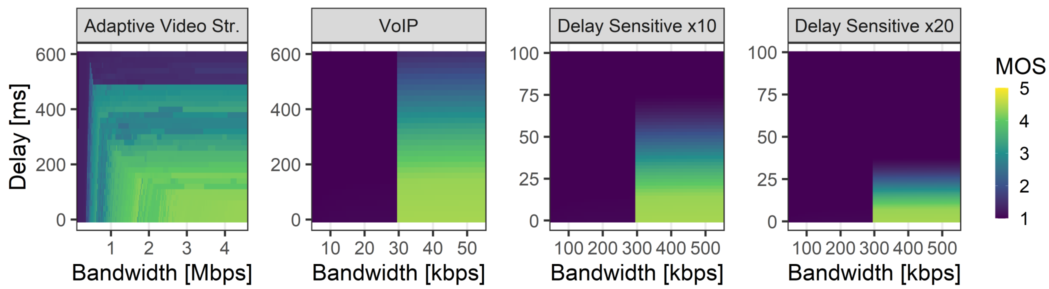

Fig. 4. QoS-to-QoE relationship for exemplary applications. While the ITU-T P.1203 model is used for adaptive video streaming, the ITU-T G.107 e-model is employed in case of VoIP and extrapolated towards delay sensitive apps.

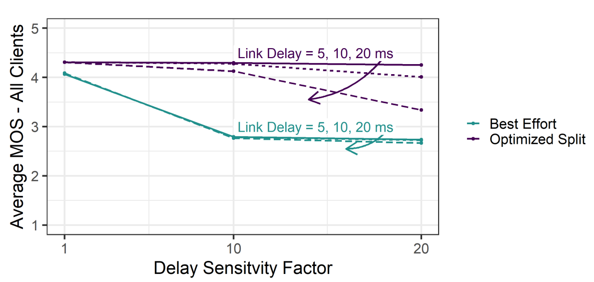

Fig. 5. Effects of different resource allocation schemes on QoE.

well as link capacity usage under different conditions. We employ a simple scenario with a topology as in [11] where application clients and servers are interconnected with a shared link with a predefined capacity and delay. In the presented case, the link has a capacity of 25 Mbps and delays of 5, 10, and 20 ms while 5 to 10 clients per application are placed in the network.

Figure 5 illustrates the development of the QoE in different simulation scenarios under two schemes: the Best Effort (BE) mode which does not perform any kind of explicit resource allocation and an *optimized capacity split* that maximizes the mean overall QoE by allocating a fraction of the available capacity to each of the applications.

For different delay sensitivity factors along the x-axis as well as different link delays and resource allocation schemes denoted by the respective line styles and colors, the y-axis shows the resulting average QoE among all application clients on the 1-to-5 Mean Opinion Score (MOS) scale. Although the BE mode is capable of coping with the requirements of current applications, i.e., a delay sensitivity factor of 1, the MOS level drops sharply to values below a MOS score of 3 as soon as the sensitivity is increased. This happens regardless of the link delay. In contrast, the optimized split can ameliorate this and maintain a good QoE even when the delay sensitivity is increased by a factor of 20 unless the baseline link delay is prohibitively high. In the latter case, the most demanding requirements cannot be met.

We envision several further options that go beyond this optimized but static split. On the one hand, dynamic re-allocation of resources according to load fluctuations can contribute to maintaining a high QoE throughout the day when users' application usage profiles shift, e.g., from work hours to evenings. On the other hand, prioritization of delay-sensitive applications can allow utilizing the link capacity to a larger extent while limiting the impact on QoE. However, both options contribute to an increase of control plane complexity and therefore require a careful analysis of trade-offs w.r.t. aspects such as scalability and performance.

V. Discussion and Conclusion

The 5G and beyond visions have changed the business and regulatory context fundamentally. When net neutrality was introduced 20 years ago, the context was that "throttling of VoIP" is not acceptable while the Internet access service threatened the main Telco business model. Today, the situation is quite opposite: Fixed and mobile broadband and the Internet access service are fundamental in the current business model. Moreover, the 5G visions are now pushing high expectations towards the telcos to deliver continued digital transformation of industry and society.

Driven by new and emerging societal needs like sustainability of ICT, this calls for a renewed look at net neutrality and puts pressure towards energy and resource efficient solutions, like the ones discussed in this article. An evolved net neutrality scheme introduces challenges from both a research and a regulatory point of view. This includes advanced monitoring of technical QoS and non-technical parameters like user fairness and non-discrimination.

Many market-oriented studies have shown the potentials of 5G and connectivity services in the various verticals. Such a demand combined with the potential of transferring lessons learned and existing building blocks from non-public connectivity services paves the way for smart Public Interconnected Networks and Services (PINS). As a result, both cost and solution uncertainties are lowered and make the situation "largely different this time".

We have discussed the unprecedented expectations upon 5G and B5G in terms of high-performing and predictable end-to-end connectivity. To achieve this in a private and local network context is not hard, however, the grand challenge is to ensure such Specialized Connectivity Services on-demand (SCS) delivered across public interconnected networks. We review the main stumbling blocks associated with this challenge. Moreover, we propose key service concepts and solution elements, as well as design directions and roadmap principles that contribute towards a holistic approach to overcome the stumbling blocks. The above makes us optimistic that a change is possible and that SCS and smart PINS might become a reality. Accompanying the analysis, a number of research questions are proposed for further work in this direction.

ACKNOWLEDGMENT

This work has been performed in the framework of EC H2020 TeraFlow (101015857).

REFERENCES

[1] M. Trevisan, D. Giordano, I. Drago *et al.*, "Five Years at the Edge: Watching Internet from the ISP Network," *IEEE/ACM Transactions on Networking*, vol. 28, no. 2, 2020.

[2] W. Lehr, D. Clark, S. Bauer *et al.*, "Whither the Public Internet?" *Journal of Information Policy*, vol. 9, pp. 1–42, 2019.

[3] C. Jacquenet, "Optimized, Automated, and Protective: An Operator's View on Future Networks," *IEEE Transactions on Network and Service Management*, vol. 18, no. 2, pp. 1350–1359, 2021.

[4] K. Kilkki and B. Finley, "In Search of Lost QoS," in *Internet Architecture Board Workshop on Measuring Network Quality for End-Users*, 2021.

[5] H. Balakrishnan, S. Banerjee, I. Cidon *et al.*, "Revitalizing the Public Internet By Making it Extensible," *ACM SIGCOMM Computer Communication Review*, vol. 51, no. 2, pp. 18–24, 2021.

[6] H. Lønsethagen, A. Vazquez-Castro, F. Gil-Castiñeira *et al.*, "Service Level Awareness and open multi-service internetworking," NetWorld 2020 ETP, White paper, 2016. [Online]. Available: http://www.networld2020.eu/wp-content/uploads/2016/07/NetWorld2020_WP_Service-Level-Awareness_Final_June-16.pdf

[7] G. Fairhurst and C. Perkins, "Considerations around Transport Header Confidentiality, Network Operations, and the Evolution of Internet Transport Protocols," RFC 9065, 2021. [Online]. Available: https://www.rfc-editor.org/info/rfc9065

[8] P. Alemany, A. Román, R. Vilalta *et al.*, "A KPI-Enabled NFV MANO Architecture for Network Slicing with QoS," *IEEE Communications Magazine*, vol. 59, no. 7, pp. 44–50, 2021.

[9] A. D. Ferguson, A. Guha, C. Liang *et al.*, "Participatory networking: An API for application control of SDNs," *ACM SIGCOMM computer communication review*, vol. 43, no. 4, pp. 327–338, 2013.

[10] European Commission, "The Smart Networks and Services Joint Undertaking," 2021. [Online]. Available: https://digital-strategy.ec.europa.eu/en/policies/smart-networks-and-services-joint-undertaking

[11] M. Bosk, M. Gajić, S. Schwarzmann *et al.*, "Using 5G QoS Mechanisms to Achieve QoE-Aware Resource Allocation," in *International Conference on Network and Service Management (CNSM)*, 2021.

[12] Q. Zhang, J. Liu, and G. Zhao, "Towards 5G Enabled Tactile Robotic Telesurgery," *arXiv preprint arXiv:1803.03586*, 2018.

BIOGRAPHIES

Håkon Lønsethagen is a senior researcher at Telenor Research. His current focus is on inter-NSP differentiated and assured services quality and business models, and management of SDN, NFV in beyond 5G architectures. He is currently part of the 6G IA Governing Board.

Stanislav Lange is a postdoc at the Department of Information Security and Communication Technology at NTNU in Trondheim, Norway. His research is focused on softwarized networks, performance analysis, and system modeling.

Thomas Zinner has been associate professor at NTNU, Norway since 2019. His research interests cover cognitive network management and network softwarization with particular focus on performance and security aspects.

Harald Øverby works as a Professor at NTNU. His interests are in the areas of: Digital Economics, Business Models, Internet Law, Regulations and Governance, and Internet Architectures. He has published over 90 papers in international and national journals and conference, and participated in several research projects.

Luis M. Contreras holds a M.Sc. on Telecoms (1997) from Universidad Politécnica of Madrid, and a M. Sc. (2010) and a Ph.D. (2021) on Telematics from Universidad Carlos III of Madrid. Since August 2011 he is part of Telefónica I+D / CTIO unit.

Nicola Ciulli is co-founder and head of Research & Innovation at Nextworks. His current activities concern architectures and protocols for network management and orchestration, with several aspects of SDN and NFV applied to transport networks and 5G/6G. He is currently sitting in the Governing Board of the 6G IA and in the Steering Board of the NetworldEurope ETP, and has (co)authored 70+ scientific publications.

Emmanuel Dotaro is the director of ICT and Security labs at Thales Secure Communications and Information Systems. He serves in both 5G-IA and ECSO PPP's associations boards as well as various Technical Committees or clusters of the digital ecosystem. His current areas of interests are network virtualization and softwarization, cloud brokering, as well as related security topics.