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 high-performing connectivity services in both public and private networks. Connectivity services that offer performance guarantees along multiple Quality of Service (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 does not exist. This especially true for interconnections across networks and support for mixed traffic modes that go beyond traditional best-effort. 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 essentially locked into the status quo. We identify the stumbling blocks and propose a set of solution elements to take us across these hurdles, alongside related research topics. This includes an approach to "Multi-Level Best-Effort (MLBE)" and suggestions for evolving net neutrality regulation. Models and simulations show how a mixed traffic mode approach provides anticipated benefits, and we provide arguments 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 B5G, 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. This illustrates an increasing demand for beyond best-effort modes for public networks [2] which are more sustainable than and can complement traditional overprovisioning [3], [4].

As a result, network operators are facing challenges in multiple directions. Dynamic resource allocation is required 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 SCS offerings 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 B5G visions are changing the business and regulatory context of future telecommunications. We now have expectations from verticals, the public sector, and governments that B5G 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 beyond today's best-effort Internet. Such an approach 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. Our work complements other efforts such as [5] and [6] which acknowledge the need for going beyond best-effort and propose tackling ossification, e.g., by means of a push towards extensibility. We provide a more holistic view by including business and regulatory perspectives, as well as concrete suggestions for traffic treatment mechanisms alongside evidence for their potential.

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. We provide a taxonomy of these challenges (Section II) including a vision towards evolving net neutrality regulation, and identify key solution elements and service concepts (Section III) that we argue are needed in order to overcome them. The overall desired outcomes are *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 Small and Medium-sized Enterprises (SMEs) that cannot build their own global backbone networks to offer Specialized Application Services (SAS) globally. In Section IV, we present illustrative numerical examples to demonstrate the expected benefits of allowing differentiated connectivity to share common resources. We discuss the appropriateness, timeliness, and feasibility of the proposed approach, and conclude the article in Section V.

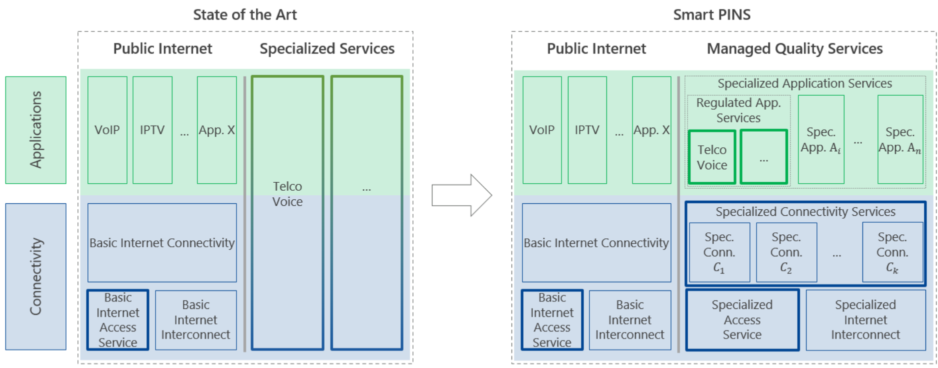

Fig. 1. Envisioned Smart PINS ecosystem with increased innovation potential and disaggregated silos. Thick highlighted boxes denote areas that are subject to regulation.

II. Stumbling Blocks and Uncertainties 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 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 Key Performance Indicators (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 Virtual Private Networks (VPNs), making operators hesitant towards entering into such new business models.

C. Privacy and Information-sharing 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 set of features or just a high level overview. 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 to B5G networks. However, it may be that

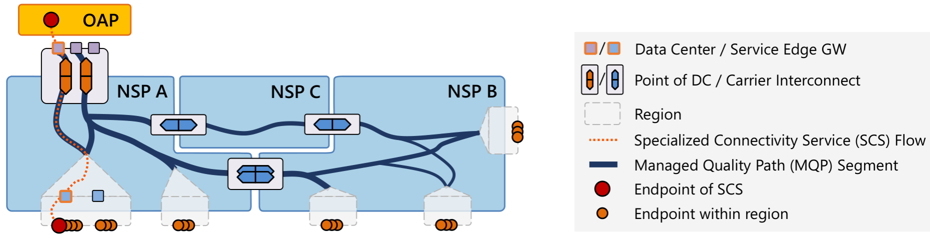

Fig. 2. Multi-domain scenario: managed quality path infrastructure with an exemplary specialized connectivity service. In this context, NSPs A and B act as edge NSPs whereas NSP C acts as transit NSP.

B5G technology offers functionality that is legally limited by current NN regulations, thus prohibiting the full exploitation of the business potentials of B5G in public networks.

We argue that NN regulations should evolve to allow SCS in an open and neutral way with equal access for all. In particular, we also emphasize the separation between connectivity and application layers as illustrated in the right-hand side of Figure 1. Thereby, the currently existing silos are opened up through a generic SCS layer that complements the current Basic Internet Access Services, thus contributing to increased flexibility and innovation potential. We anticipate that regulated and unregulated application services will persist, both on top of SCS.

III. Service Concepts, Main Solution Elements, and Recommendations

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. 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.

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 reflecting this heterogeneity are required for efficient traffic handling. Depending on the tolerable amount of complexity, particularly control plane complexity, 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 even be re-routed via longer paths. Evolving from the "best-effort" traffic mode of today's Internet as starting point, we propose and discuss four main traffic modes. These include three best-effort modes that differ from each other, allowing 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 by mechanisms like Weighted Fair Queueing (WFQ) or routing via a shorter path. 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. Examples for applications using the BG, BQ, and IQ mode respectively could be the automatic download of an Operating System (OS) update, an on-demand video stream, and a live video stream which have increasingly strict QoS requirements while not necessarily being mission-critical and requiring strict performance guarantees. 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 requires network performance that is significantly higher or more stable than what the IQ mode offers. This will require mechanisms that are more complex than those of the IQ mode. That is, to enable the AQ mode, the QoS, resource allocation, 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 generally 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 and which will enable managed traffic aggregates to reach Service Edge Gateways (SEGs).

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 SEGs or interconnection points need to be provisioned whereas the remainder can be carried by

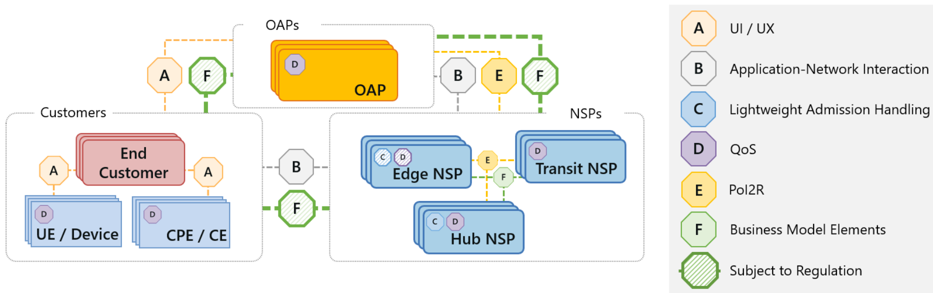

Fig. 3. Overview of solution elements, interfaces, as well as regulatory considerations between and within the main stakeholders, i.e., users, application providers, and network providers. Depending on the context, NSPs can play the roles of edge, transit, or hub NSPs which come with their respective internal mechanisms and regulatory conditions.

a suitable, already provisioned, and well-dimensioned MQP. These pre-established MQPs 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 session, different traffic modes might be employed and aligned at the MQP and the SCS session level. Figure 2 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. The core idea of the MQP service is to enable dynamic traffic engineering, intelligent management, and configuration of coarse traffic aggregates with their services. Additionally, it may support remote peering. From this perspective, current peering and transit services need to be evolved to provide *i)* more sophisticated SLAs, and *ii)* more fine-grained regions, both in a spatial / geographical and technological sense.

Since SCS sessions use MQPs, only paths from end-users and application servers to local GW endpoints matching the corresponding MQP need to be provisioned at either end of the end-to-end SCS session. We refer to these endpoints as Data Center Gateway (DC GW) and 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. A core challenge in this context will be to strike a balance between covering SCS needs and scalability, e.g., by merely checking policies at the gateways. This way, signaling and QoS handling could 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 research challenges for each of the elements. 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, with an emphasis on the interfaces between and within the customer, application, and network domains, is provided in Figure 3.

User Interaction / Interface (UI) / User Experience (UX): UI / UX aims at enabling the customer to decide the appropriate service level. This can be done either in an explicit manner, e.g., via UI and UX dialogues or implicitly based on relevant characteristics of the environment and end devices in use. Research challenges involve appropriate means of user guidance and interaction, and establishing consistent approaches for expressing service level expectations and indicators.

Application-Network Interaction (ANI): ANI aims at interfaces to allow requests for SCS and corresponding NSP offerings. A potential realization can be similar to application-initiated reservation strategies as proposed for Software-Defined Networking (SDN) in the context of participatory networking [9], or to IETF efforts like Network Service Headers (NSH) and Application-Aware Networking (APN). The NSP could provide templates expressing possible QoS value ranges - e.g., target and lower bound - in terms of supported throughput, latency, and packet loss. Challenges include API definitions, information and data models, how to ensure scalability, portability, and efficiency for the application developer, and dynamic re-negotiations taking application elasticity into consideration.

Lightweight and Class-based Admission Handling (LAH): The main objective of LAH is to ensure that the volume of admitted SCS sessions will not make the BQ mode and general quality level suffer beyond specific committed performance levels. For that, the NSP needs to monitor application traffic rate and behavior per OAP and perform class-based and per-OAP policing by various means. Challenges include planning of addresses and identifiers, understanding the impact of the available policing and admission mechanisms within NSP, and the design of scalable admission handling solutions, e.g., on a per-class or SCS level.

QoS: QoS mechanisms allow (re-)aligning the allocation of available resources with users' requirements. They are available at multiple locations within the User Equipment (UE) / device, edge NSP, and OAP domains. While today's mechanisms mainly focus on bandwidth allocation, advances from the areas of deterministic, time-sensitive, and high-

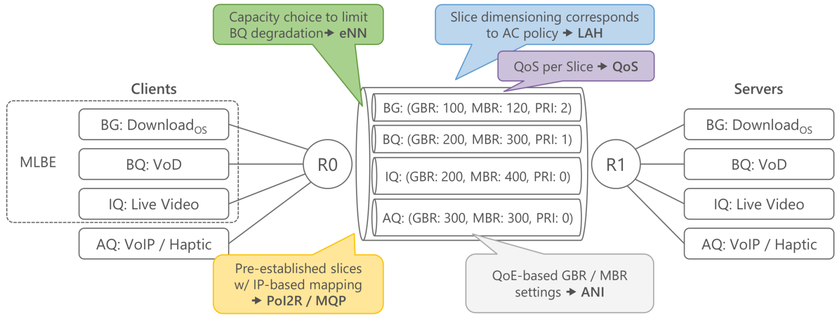

Fig. 4. Simulation scenario for evaluating the proposed MLBE approach and coverage of corresponding key solution elements.

precision networking can be leveraged to address delay requirements. Challenges include the choice of simple but sufficient QoS mechanisms that account for technical and economical constraints, finding the appropriate configuration of a specific QoS mechanism relative to the traffic mode, and integrating appropriate mechanisms to reduce network delays.

Point-of-Interconnect-to-Region (PoI2R) Service: The PoI2R is an interconnection service, where the notion of "region" can be in a spatial / geographical or technological sense, e.g., a range of IP prefixes, and which provides the possibility to assign certain performance criteria like throughput and delays. Design and implementation details may differ depending on the specific type of peering, i.e., direct peering between edge NSPs like eyeballs and OAPs, or peering with transit NSPs. Further unsolved challenges in this context include security, appropriate information and data models - e.g., to augment IP prefixes with additional attributes - and how to address potential fairness and neutrality issues.

Business Model Elements (BME): Along with the ANI and PoI2R service variants and enablers, new BME and charging principles will be required. That includes also hybrid money flows from the customer to both the OAP and NSP, and to optionally support Initiating Party Network Pays (IPNP) for two-way connectivity across NSPs. Relevant challenges include an understanding of the appropriate business models for the different stakeholders and assuring compatible incentives to enable a healthy market.

Net Neutrality: New interfaces and information flows will enable the alignment of traffic modes with application requirements, thereby challenging the current scope of net neutrality, which is purely related to technical parameters. A key question in this context revolves around expressing and measuring the general performance level of the BQ mode to ensure protection of the basic IAS.

C. Regulation and Evolution Towards Smart PINS

An evolutionary approach to Smart PINS is needed demonstrating first that costs can be managed well, and that the effects are achievable at the local level.

Trials and pilots will be important to provide evidence of the capabilities and measurement techniques that can enable the anticipated solution elements. With growing maturity of the solution elements, pre-commercial or commercial pilots can enter the market. In the early commercial phase, the existing well-dimensioned best-effort IP peering will be sufficient for many cases, while the MQP solution elements can be introduced subsequently. In particular, we consider ANI, LAH, and BME to be critical already in the early stages.

Economically, the business models of the OAP and NSP need to evolve, both between the different providers and towards the end-users. This will start with business models between an OAP reaching customers within a directly connected NSP, and then evolve to models enabling to reach end-users within remote NSPs. The hub NSP role will be important for scalability, allowing an edge NSP to reach numerous remote edge NSPs. In parallel, pricing models towards end-users will evolve to ensure a correspondence between the price paid and the resources used, hence incentivizing responsible and ecologically viable resource usage.

Finally, we anticipate the need for evolving the NN regulation to ensure that open and equal access to SCS can be supported at a global scale. This will be critical in order to realize the vision of smart PINS and to achieve the desired scalability of MLBE traffic modes.

IV. Evidence of Potential

In order to assess the potential of the proposed MLBE approach, we design a simulation that covers the main solution elements outlined in the previous section and compare its performance characteristics to those of a traditional best-effort approach. The simulation scenario is illustrated in Figure 4 and includes a heterogeneous mix of application servers and clients corresponding to four traffic modes: OS-initiated file downloads (DL) in the BG, user-initiated video streaming (VoD) with BQ, live video streaming (LVD) with IQ, and a highly delay-sensitive application emulating the exchange of haptic feedback (HAP) that requires AQ. The applications differ not only in terms of their requirements, but also w.r.t. their elasticity and transport protocols. For instance, the video applications use adaptive streaming over HTTP and can adjust to given network conditions whereas DL and HAP applications rely on their respective TCP and UDP connections.

Leveraging the methodology from [10], the 1 Gbps link between servers and applications is sliced using a Hierarchical Token Bucket (HTB) scheduler, which enables us to include the following solution elements:

- Per-slice Guaranteed and Maximum Bit Rate (GBR, MBR) settings as well as priorities allow for a QoS-based

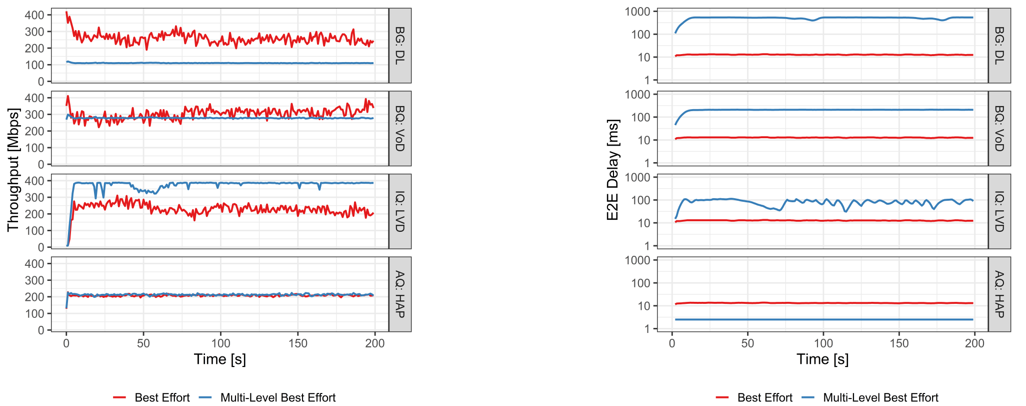

(a) Time series of aggregated per-application throughput.

(b) Time series of aggregated per-application delays.

Fig. 5. Impact of different traffic handling strategies on the throughput and delay performance of a heterogeneous set of applications.

resource allocation. Since these settings are on coarse per-application and IP address range granularity, they can be provisioned in advance as per the PoI2R / MQP concepts.

- Given the slice parameters, the static number of clients per application corresponds to an admission control policy which limits those numbers.
- Slice parameters are chosen in a QoE-aware manner and are tailored to applications' requirements, hence covering aspects of ANI.
- By carefully dimensioning the resources for the BQ mode, we ensure that the user-perceived application quality remains unchanged. This behavior aligns well with the goals related to evolving net neutrality regulations.

In particular, the GBR, MBR, and priority settings are chosen to represent the typical characteristics and requirements of the traffic modes: an isolated high-priority AQ mode as well as BG, BQ, and IQ modes with increasing levels of priority and potential for intra-slice capacity borrowing. For the following study, 110 active clients for each of the BG, BQ, and IQ modes and 825 clients of the AQ mode are simulated. The number of the latter is higher to compensate for the lower capacity consumption per client.

The graphs in Figure 5 show the aggregated per-application throughput and delay characteristics over the course of the simulation. Each chart features two curves corresponding to the behavior in a Best-Effort (BE) regime without differentiation and the proposed MLBE approach. We make several observations on the throughput performance displayed in Figure 5a: highly fluctuating throughput values for the three TCP-based applications in the BE case are caused by the heterogeneous application mix and the clients' dynamic behavior which prevents an equilibrium with a fully fair share. In contrast, the homogeneity of applications within each class in the MLBE case leads to significantly more stable throughput values which are constrained by the set GBR and MBR values. Being UDP-based, the HAP application clients manage to send all their packets and achieve the same throughput in both the BE and MLBE conditions. Hence, the smoothing effect observed for other applications does not occur in the case of HAP. Please further note that minor throughput fluctuations in HAP traffic are caused by application-level behavior and happen independently of the specific traffic treatment.

However, the delay performance reported at the bottom of Figure 5b highlights that using BE results in delays above 12 ms, which are outside the 3-10 ms range reported as a requirement for haptic applications in [11]. Without differentiation, all applications experience the same delay in case of BE. When using MLBE on the other hand, the delays of the HAP clients decrease to an average of 2.5 ms, but come at the price of increased delays for all other applications. Since these are more delay-tolerant, however, the user-perceived application quality is not negatively impacted in the BQ and IQ case.

To quantify the QoE for video streaming applications on the 1-to-5 Mean Opinion Score (MOS) scale, we leverage the ITU-T P.1203 model [12]. Our results show that while the MOS for the VoD application in the BQ mode remains unchanged at 4.18 in both the BE and MLBE cases, the IQ mode's MOS drastically improves from 2.83 to 4.38 when following the MLBE approach. This is mainly due to the increased amount of link capacity that is allocated to the IQ traffic mode.

Note that although the delay experienced by the background download increases substantially and the lowered throughput causes longer download times, it maintains a stable bit rate throughout the experiment and is never starved. Since we assume the download to be an OS-initiated background process, we do not expect the aforementioned effects to have a noticeable effect on user experience.

In summary, we have illustrated the potential of the combined MLBE and AQ approach to support emerging highly demanding applications and to selectively improve the performance of existing applications, while limiting the impact on background and basic traffic modes in a way that preserves user-perceived application quality.

V. Discussion and Conclusion

The B5G 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's situation is quite opposite. Fixed, mobile broadband, and Internet access services are fundamental in the current business model. Moreover, B5G visions are pushing high expectations towards telcos to deliver the ongoing digital transformation of industry and society.

Driven by emerging societal needs, such as "green ICT" and the sustainability of ICT, there is a call for a renewed look at net neutrality. This also puts pressure towards energy and resource efficient solutions, like those discussed in this article. Furthermore, evolving the notion of net neutrality introduces research-related and regulatory challenges, including 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 PINS. As a result, both cost and solution uncertainties are lowered and make the situation "largely different this time".

Accompanying an analysis of the ecosystem, this article provides key service concepts and related research questions required for the evolution towards smart PINS.

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, pp. 1–12, 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, Accessed on 2023-08-23. [Online]. Available: https://www.iab.org/wp-content/IAB-uploads/2021/09/Kilkki-In-Search-of-Lost-QoS.pdf

[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] 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, pp. 283–291.

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

[12] I.-T. S. G. 12, "ITU-T Recommendation P.1203: Parametric Bitstream-based Quality Assessment of Progressive Download and Adaptive Audiovisual Streaming Services over Reliable Transport," International Telecommunication Union, Tech. Rep., 2017.

Biographies

Håkon Lønsethagen is a Senior Research Scientist at Telenor Research. Since 1990 he has been working with telecom network and service management, orchestration, and control, including systems integration and distributed systems frameworks, architectures, and middleware. Over the last years his activities have addressed inter-provider network services and business models, Internet evolution, SDN, NFV and 5G technologies and ecosystems. He is currently part of the 6G IA Governing Board.

Stanislav Lange is an associate professor 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 received his M.Sc. and PhD degrees in Computer Science from University of Würzburg in 2007 and 2012. Prior to joining NTNU in 2019 he was lead research groups at TU Berlin and University of Würzburg. He joined NTNU in 2019 and has been full professor since 2022. His research interests cover cognitive network management and network softwarization with particular focus on performance and security aspects.

Harald Øverby has a M.Sc. in Computer Science (2002) and a PhD in Telematics (2005) from Norwegian University of Science and Technology (NTNU). Since 2022, Øverby is the Provost for Student Learning at BI Norwegian Business School. Prior to joining BI, he was affiliated with NTNU for more than 20 years as a Head of Department and Professor. Øverby has published more than 100 academic works, covering areas including digital economics, optical networking, Internet architectures, network coding, and cybersecurity.

Luis M. Contreras holds an M.Sc. in Telecommunications from UPM (1997), an M. Sc. in Telematics jointly from UC3M and UPC (2010), and a Ph.D. cum laude in Telematics from UC3M (2021). Since August 2011 he is part of Telefonica working on scalable networks and their interaction with cloud and distributed services, evolution to B5G/6G, etc. He is also an active contributor to different SDOs, such as IETF, O-RAN, or ETSI.

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.