

# Incentive-Based Traffic Management and QoS Measurements

Szilveszter Nádas\*, Balázs Varga\*, Luis M. Contreras<sup>‡</sup>, Sándor Laki\*\*

\*Traffic Analysis and Network Performance Laboratory, Ericsson Research, Budapest, Hungary

<sup>‡</sup>Transport Technology and Planning, Telefonica CTIO, Madrid, Spain

\*\*Communication Networks Laboratory, ELTE Eötvös Loránd University, Budapest, Hungary

{szilveszter.nadas,balazs.a.varga}@ericsson.com, luismiguel.contrerasmurillo@telefonica.com, lakis@inf.elte.hu

**Abstract**—We argue that the current Best Effort Quality of Service (QoS) strategy of the Internet is one of the reasons why meaningful quality indicators and measurements are hard to define over the Internet. At the same time, Session-based QoS has emerged for mobile communication, VoIP and IPTV services. However, it is way too complex for widespread usage and cannot out-weight the simplicity being the primary advantage of Internet’s success. Incentive-based strategies have been proposed as lightweight alternatives to traditional solutions in the past decade. As an example, we explore the Core Stateless QoS strategy and show that it is simple, it still provides a fine control of QoS, and that it can also provide a network wide congestion metric. It enables inter-domain policy translation by implementing hierarchical resource sharing policies of arbitrary depth. Furthermore, it also enables network operators and subscribers to carry out meaningful measurements, ensure and verify Service Level Agreements (SLAs) and support network-endpoints signals. We propose future work in this area, including the integration of measurements into the Traffic Management strategies, that are needed for the widespread deployment.

**Index Terms**—Traffic Management, QoS, Incentive, Measurement, Core-stateless, SLA

## I. INTRODUCTION

Despite extensive research and standardization in the area of Quality of Service (QoS), most of the developed solutions have not been deployed in practice [1], [2]. Proponents of overprovisioning argue that it is much easier and more efficient to add capacity when needed than to build and maintain complex QoS mechanisms that only provide minor improvement during congestion [1]. Network congestion together with the ways of avoiding its impacts on the end users is a recurrent aspect along time, especially nowadays where the network is becoming a critical societal asset. Recently, the unfortunate pandemic situation due to COVID-19 made evident this critical fact, revealing the need of defining proper mechanisms for improving its robustness and ensuring sufficient quality of experience for end users.

The report in [3], based on measurements from RIPE Atlas probes, states that last-mile network congestion bottlenecks grew from a usual 10% of AS showing congestion up to 55% of them experiencing issues. Interestingly, also that report mentions situations where operators considered speed upgrades for certain (less-favored) users to ensure certain levels of service, or even discussions among operators and content providers for coordinating actions in order to prevent

generalized congestion events. Some very remarkable statements provided as recommendations refer to the “... need to ensure that broadband is available to all and that Internet services equally serve different groups” and to the “... ability for networks and applications to collaborate better”.

In addition, the lack of QoS over the Internet resulted in the emergence of private platforms to support enhanced services [2]. For example most Over-The-Top providers have their own proprietary Content Delivery Network (CDN) creating a proprietary international core network (i.e., private WAN) extended by cache servers deployed close to the users. The authors of [1] argue that a paradigm change from requirements to incentives might be required, it is not clear however how these incentives would look like. We define *incentive* as a piece of information that can be used for making decisions about how to satisfy QoS. We also require it to be more lightweight than session parameters.

While some argue that the aggregation links of Access Aggregation Networks (AAN) – typically controlled by the ISPs – are not bottlenecks, with higher access speeds for the end-users they can easily become ones. To provide fair sharing of resources among users, Hierarchical QoS (HQoS) is used in some parts of the AAN, though this does not scale for the whole AAN. In the parts of the AAN where HQoS is not used, the resource sharing is controlled by the interaction of TCP Congestion Control Algorithms (CCAs), which can be very unfair not only because of the unfairness of various CCAs [4], but also because some users might use much more TCP flows. The Internet Core is even harder to manage, and most of the traffic is served by CDNs located close to the AAN anyway, the authors of [5] formalize the role of these Service Nodes (e.g., CDN caches).

The Quality of Experience of Internet services is usually very good, so a single short quality measurement is unlikely to reveal any problems. Even if a measurement is performed in a problematic state of the network it is very hard to identify the cause of the problem, partly because of the very simple QoS architecture and simplistic associated SLAs. In this paper, we focus on congestion related problems, when the service degradation is the consequence of other traffic also using a shared bottleneck. For example the traffic of a heavy user may also use a shared bottleneck; or another traffic flow of the same subscriber may cause issues. Even in a lightly loaded network




Fig. 1. Traffic Management

the shared buffers may still be filled to a high degree from time to time causing issues for latency sensitive applications also using that buffer.

## II. TRAFFIC MANAGEMENT

Fig. 1 shows *mechanisms* controlling congestion on different timescales (based on [6]). Notice that mechanisms that are good to control short term congestion are not good for long-term one and vice-versa. For example, the bufferbloat problem [7], unnecessarily high queueing delay over some Internet bottlenecks is best solved by novel AQM and Scheduling algorithms. It would be very hard to find an Internet-wide solution among other mechanisms. E.g., replacing all Congestion Controls to delay-based ones seems hardly possible. For each mechanism, there is a number of alternative *algorithms*, e.g., Cubic, DCTCP or BBR can be used as Congestion Control Algorithms. We define *strategies* as harmonized sets of algorithms, where one (or more, or zero) algorithm is used for each mechanism. Even if the right mechanism is chosen to solve a problem, it is only possible to achieve limited impact by updating a single algorithm. The update of the whole strategy is more likely to have impact. E.g., DCTCP congestion control was impossible to deploy over the Internet, but by introducing a new packet-level signal and new AQM and Scheduling algorithms, it may become deployable.

As an example the “Best Effort Internet access” strategy has unlimited resource use as Usage Policy, with a peak rate and often with a much smaller guaranteed bitrate. Overprovisioning is used as Network Dimensioning, there is no Admission Control over the Internet. Some services (e.g., Youtube) use content adaptation, and most traffic is Congestion Controlled using a TCP CCA. The CCAs are also responsible for how resources are shared among flows, though some access networks are using Hierarchical QoS or air interface scheduling to control resource sharing among users and in some cases among applications. Typically a simple FIFO AQM is used, though some cases one of the novel AQM techniques, sometimes coupled with a flow scheduler are utilized, e.g., fq-codel.

Other networking technologies included more complex algorithms, like those needed to realize the “session-based QoS” usage policy. That has its own problems: too detailed QoS requirements which are hard to provide and verify; and that almost no service actually needs (or can pay for) such hard guarantees. The reasons behind the success of the Internet

include its flat pricing and simple Traffic Management, which would not be possible if generic session-based QoS would have been introduced.

The question arises: Is it possible to create a strategy which keeps the simplicity of the “Best Effort Internet access”, but contains a somewhat richer SLA, helping in providing the right QoS for applications and in creating a meaningful measurement to verify that SLA?

In the next section we investigate few such promising proposals. We believe that the paradigm change from requirements to incentives as proposed in [1] is a way worth considering for a better and more customizable Internet service.

## III. INCENTIVE-BASED APPROACH

We define *incentive* as a piece of information, which can be used for making decisions about QoS. We also require it to be more lightweight than session parameters both in the complexity of traffic management solutions to be applied as well as in the deployment costs. As opposed to “Best Effort Internet QoS” an “Incentive-based QoS” strategy shall be based on, e.g., a more detailed SLA than peak rate and guaranteed rate; a more detailed signals from one or more of the endpoints; or request-response communication between a network element and an endpoint. In this section, we list examples for each. [8] summarizes relevant considerations on Application - Network Collaboration Using Path Signals. [9] lists potential properties that may be exposed by the network to applications.

### A. More detailed SLA

An example for more detailed SLA is to introduce monthly cap for the Broadband traffic, but exclude traffic in underloaded periods. This creates the Incentive to schedule one’s not urgent downloads for the underloaded periods, therefore decreasing the load in the busy hour.

Another such SLA is Multi-timescale Bandwidth Profile [10], which allows temporal throughput bursts for sources with “good transmission history”. By rewarding silent periods with improved performance, it provides an incentive to not overuse the network in periods when QoS is important for the user.

### B. Packet-level signals from endpoints

Many different proposals for packet-level incentives has emerged in the literature in the past decade. LoLa adds information to the packet header on whether the given packet requires low-latency or not. L4S extends that by also stating Congestion Control behaviour. Similarly, DiffServ defines per-hop behaviors (PHBs) encoded into the DSCP field of IP packets. PHBs determine how the packet shall be handled by the routers.

In addition to traditional approaches few solutions using packet-level incentives have emerged in the past decade. These solutions require extra fields in the packets where the incentive values can be encoded, while packet scheduling and drop decision solely rely on these carried values. The key advantage

of such solutions is that network nodes can operate in a flow-unaware fashion. On the other hand, incentives need to be assigned at some point of the network. In this way, the role of QoS management is shared among different network entities. In case of core-stateless resource sharing approaches [11]–[13], the packets are marked with a value calculated from the sending rate of the traffic aggregate they belong to. Then such values can solely be used in the traffic management engine to decide which packet to drop in case of congestion, ensuring weighted fairness among traffic aggregates at flow, user or application-levels.

[14] defines Qualitative Communication Service where some payload portions may be more important to applications than others. The qualitative networking approach exploits this fact by allowing senders to group payload within a packet by relative priority, then allowing the network to selectively discard portions of lesser priority when needed.

### C. Packet-level signals from the network

The network may also send packet level signals. ECN Congestion Experienced provides congestion signal without dropping packets. A proposal to provide throughput guidance to congestion control is proposed in [15].

### D. Request-Response Communication

An example for Request-Response Communication is to temporarily boost Broadband service by purchasing extra monthly cap or by upgrading the subscriber policy class for, e.g., a day. While most Internet use cases may be served without utilizing explicit communication, some strategies may still benefit from it.

## IV. EXAMPLE - CORE STATELESS TRAFFIC MANAGEMENT

We take the example of Per Packet Value (PPV) based resource sharing [13] to demonstrate different algorithms of an “Incentive-Based Core Stateless QoS” strategy.

### A. Resource Sharing Control

The PPV framework encodes Resource Sharing Policies to a Packet Value marked on each packet. Resource sharing policies are expressed by Throughput Value Functions (TVFs). Each TVF is used to label packets of a traffic aggregate where the packet value expresses the gain that is only realized if the packet is delivered (marginal utility in other words). By

applying the TVF at the right aggregate, e.g. subscriber level, resource sharing becomes independent of the number of flows by a subscriber.

Fig. 2 illustrates how the TVFs and packet values (PVs) can be used to share the bottleneck capacity between various flows. In the first case, the bottleneck capacity is 10Mbps shared between three flows. The red, blue and green curves on the right side represent the TVFs of Flow1, Flow2 and Flow3, resp. The gray dotted line illustrates the cutoff value that results in a resource allocation 0, 6.25 and 3.75Mbps for Flow1, Flow2 and Flow3, resp. This allocation is ensured by only transmitting packets with PV above the cutoff level. One can observe that Flow1 has no packet with PV above this threshold and thus it cannot even transmit a single packet. We call this cutoff value as Congestion Threshold Value (CTV). In the case of a 45Mbps bottleneck, the resulting CTV is much smaller and thus all three flows have non-zero assigned throughput. The purple dotted line represents the CTV, leading to 10, 17.5 and 17.5Mbps throughput allocation for Flow1, Flow2 and Flow3, resp. In this case, only packets below the CTV shown by the purple line are dropped (or marked with ECN CE).

In addition to the packet value, a Delay Class may be marked on each packet which might be used similarly to the L4S bit.

### B. AQM and Scheduling

A PPV capable AQM and Scheduling algorithm aims to maximize the transmitted total Packet Value, while also taking into account the delay class. By this maximization the desired resource sharing is realized. At the bottlenecks no flow identification or policy knowledge is required. When maximizing the transmitted Packet Value, the Congestion Threshold Value (CTV) emerges, where packet with value higher than the CTV are transmitted without loss. The CTV is unique at each bottleneck and can be used as a rich congestion measure that determines the allowed throughput of any affected flow or traffic aggregate. Knowing the CTVs along a network path enables the end-users to check their expected throughput-related SLAs.

The delay class is orthogonal to the Packet Value, which results in meeting resource sharing targets even among flows of different delay classes allowing, e.g., low resource priority for low delay traffic.

### C. Congestion Control

Existing TCP CCs may be incompatible with each other [4]. There is no restriction on the applied CC in the PPV system.

### D. Network Dimensioning

When the TVFs used by subscribers and the traffic dynamics are known, the network can be dimensioned for various use cases. In a *busy hour scenario* a CTV may be targeted to meet, which also defines the throughput reached by a customer with a given TVF. Other scenarios like *very high load* and *worst case load* may be defined with different CTV targets (and different emerging throughput values).




Fig. 2. Resource sharing with the PPV framework

### E. Usage Policy

The share of a subscriber is defined by its TVF. Multi-Timescale sharing can also be encoded into the Packet Values by carrying out throughput measurements on multiple timescales and then applying Multi-Timescale TVFs [10] to determine Packet Values.

### F. Inter-Domain Policy Translation

A unique capability of the PPV framework is that hierarchical resource sharing policies of arbitrary depth can be encoded to the (series) of single PV. E.g. if subscribers can mark their own packets they can encode the sharing among their flows into such a TVF [16]. Also on domain border a packet marker [17] can mark packets according to the TVF of the aggregate, while keeping existing policies of the aggregate, thereby adding a new layer to hierarchical resource sharing, without explicit information about existing layers. This rich hierarchical policy translation is hard to achieve with other solutions, e.g. by DSCP PHBs.

### G. Measurements

By formalizing resource sharing the end-point may utilize richer measurements. They can measure the CTV provided by the network and the delay experienced by the different delay classes. It is possible to measure how CTV changes during the day, thereby allowing insights on how the load in the network changes.

The network may also feed-back CTV to the endpoints to help in Congestion Control and Content Adaptation. As the CTV is very easy to verify with end-user measurements, there is little incentive for the network to falsify it. For critical periods a premium service might be possible to buy, implemented by a modified TVF resulting in a higher amount of resource shares.

## V. CONCLUSION

More advanced QoS is hard to define without actually updating Usage Policy for the end users. We gave an overview of traffic management strategies and argued that meaningful QoS measurements should be part of the strategy itself. That way the defined measurements could be supported by Traffic Management algorithms including the Usage Policy. We have shown that the Core Stateless Traffic Management is a promising strategy due to its richness, simplicity and because it includes a rich congestion measure, which can harmonize different traffic management algorithms including end-user QoS measurements.

## VI. FUTURE WORK

We propose discussion and further research in the area of Incentive-Based Traffic Management. Some possible tasks are as follows: 1) Investigating how to create QoS related incentives, along which actors in the Internet can cooperate to achieve better QoS, and where misuse is discouraged. 2) Keeping pricing simple while enabling new services and higher, consistent service availability. 3) Understanding the

importance of Traffic Management strategies to provide better QoS. 4) Creating strategies where richer user measurement is possible. This likely requires richer SLAs, and can be supported by the network sending signals about its congestion state. 5) Investigating how to design Traffic Management algorithms that support measurements. 6) Investigating how a Traffic Management strategy can be best supported by novel SLAs and how these SLAs can be marketed towards end-users.

## REFERENCES

- [1] K. Kilkki and B. Finley, "In Search of Lost QoS," *arXiv preprint arXiv:1901.06867*, 2019.
- [2] K. Claffy and D. D. Clark, "Adding enhanced services to the internet: Lessons from history," *Journal of Information Policy*, vol. 6, no. 1, pp. 206–251, 2016.
- [3] J. Arkko, S. Farrell, M. Kühlewind, and C. Perkins, "Report from the IAB COVID-19 Network Impacts Workshop 2020," RFC 9075, Jul. 2021. [Online]. Available: <https://rfc-editor.org/rfc/rfc9075.txt>
- [4] F. Fejes, G. Gombos, S. Laki, and S. Nádas, "Who will save the internet from the congestion control revolution?" in *Proceedings of the 2019 Workshop on Buffer Sizing*, 2019, pp. 1–6.
- [5] H. Balakrishnan, S. Banerjee, I. Cidon, D. Culler, D. Estrin, E. Katz-Bassett, A. Krishnamurthy, M. McCauley, N. McKeown, A. Panda *et al.*, "Revitalizing the public internet by making it extensible," *ACM SIGCOMM Computer Communication Review*, vol. 51, no. 2, pp. 18–24, 2021.
- [6] R. Jain, "Congestion control and traffic management in ATM networks: Recent advances and a survey," *Computer Networks and ISDN systems*, vol. 28, no. 13, pp. 1723–1738, 1996.
- [7] J. Gettys, "Bufferbloat: Dark buffers in the internet," *IEEE Internet Computing*, vol. 15, no. 3, pp. 96–96, 2011.
- [8] J. Arkko, T. Hardie, and T. Pauly, "Considerations on Application - Network Collaboration Using Path Signals," Internet Engineering Task Force, Internet-Draft draft-arkko-iab-path-signals-collaboration-00, Jul. 2021, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-arkko-iab-path-signals-collaboration-00>
- [9] T. Enghardt and C. Krähenbühl, "A Vocabulary of Path Properties," Internet Engineering Task Force, Internet-Draft draft-irtf-panrg-path-properties-03, Jul. 2021, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-irtf-panrg-path-properties-03>
- [10] S. Nádas, B. Varga, I. Horváth, and Mészáros, "Bandwidth profile for multi-timescale fairness," in *2020 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2020, pp. 1–8.
- [11] Z. Yu, J. Wu, V. Braverman, I. Stoica, and X. Jin, "Twenty Years After: Hierarchical Core-Stateless Fair Queueing," in *NSDI*, 2021, pp. 29–45.
- [12] M. Menth and N. Zeitler, "Fair resource sharing for stateless-core packet-switched networks with prioritization," *IEEE Access*, vol. 6, pp. 42 720–42 720, 2018.
- [13] S. Laki, S. Nádas, G. Gombos, F. Fejes, P. Hudoba, Z. Turányi, Z. Kiss, and C. Keszei, "Core-Stateless Forwarding With QoS Revisited: Decoupling Delay and Bandwidth Requirements," *IEEE/ACM Transactions on Networking*, vol. 29, no. 2, pp. 503–516, 2020.
- [14] I. FG-NET2030, "New services and capabilities for network 2030: description, technical gap and performance target analysis," *FG-NET2030 document NET2030-O-027*, 2019.
- [15] A. Jain, A. Terzis, H. Flinck, N. Sprecher, S. Arunachalam, K. Smith, V. Devarapalli, and R. B. Yanai, "Mobile Throughput Guidance Inband Signaling Protocol," Internet Engineering Task Force, Internet-Draft draft-flinck-mobile-throughput-guidance-04, Mar. 2017, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-flinck-mobile-throughput-guidance-04>
- [16] F. Fejes, S. Nádas, G. Gombos, and S. Laki, "A Core-Stateless L4S Scheduler for P4-enabled hardware switches with emulated HQoS," in *IEEE INFOCOM 2021-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2021, pp. 1–2.
- [17] S. Nádas, Z. Turányi, G. Gombos, and S. Laki, "Stateless resource sharing in networks with multi-layer virtualization," in *ICC 2019-2019 IEEE International Conference on Communications (ICC)*. IEEE, 2019, pp. 1–7.