

MR-452.4

QED Uses in Lab Evaluation & Network Design

Issue: 3
Issue Date: February 2021

Issue History

Issue Number	Approval Date	Release Date	Issue Editor	Changes
1	9 February 2021	9 February 2021	Gavin Young (Vodafone)	Original.

Comments or questions about this Broadband Forum Marketing Report should be directed to info@broadband-forum.org.

Editor Gavin Young (Vodafone)

**Work Area Director(s),
Access and Transport
Architecture (ATA) Work
Area** David Sinicrope (Ericsson)

Contributors Peter Thompson (PNSol)
Neil Davies (PNSol)
Ken Kerpez (ASSIA)

Table of Contents

Executive Summary..... 4

Introduction 5

1 References and Abbreviations..... 6

 1.1 References..... 6

 1.2 Abbreviations..... 6

2 Applicability of ΔQ to Lab Emulation 7

 2.1 Quality Attenuation Emulator 7

 2.2 Lab Emulation Environment..... 7

 2.3 A/B testing..... 8

3 ΔQ Design and Lab Evaluation Use cases 10

 3.1 De-risking Application Deployment..... 10

 3.2 Optimizing Speed Tests..... 10

 3.3 Measuring Cost/Benefit of Multi-access Edge Computing (MEC)..... 10

 3.4 Physical Equipment Performance Benchmarking for Selection..... 10

 3.5 Virtual Network Function (VNF) Evaluation & Abstraction..... 11

 3.6 Equipment Configuration Impact..... 12

 3.6.1 *Comparing GPON Configuration Impact*..... 12

 3.6.2 *Impact of Turning on QoS on GPON*..... 13

 3.7 Input to Network Architecture Design/Analysis 14

 3.8 Use of ΔQ in Systems Integration 15

4 Summary and Conclusions..... 16

List of Tables

Table 1: Equipment interface speed and resulting delay to bits & packets..... 11

List of Figures

Figure 1: Lab emulation 8

Figure 2: A/B testing..... 9

Figure 3: Example of abstraction of ΔQ observable outcomes..... 12

Figure 4: Comparing GPON Configuration Impact..... 13

Figure 5: Impact of turning on QoS on GPON..... 14

Figure 6: Potential roles of Quality Attenuation measurements in the technology lifecycle 16

Executive Summary

This Marketing Report contains use-cases for using the Delta Q framework within a network architecture/design and Lab evaluation environment.

Use of performance measurement and monitoring tools is commonplace in the “In life” operations phase of the broadband network technology lifecycle. Network Operations Center (NOC) and field operations teams have a range of tools and testers to leverage for this purpose. Many of these tools (such as TWAMP probes) could be enhanced to provide Quality Attenuation measurements [1] that will give greater insight into network health and assist root cause analysis.

However, Quality Attenuation for performance measurement and analysis can also play an extremely valuable role in other phases of the technology lifecycle. These include design and lab evaluation of new equipment, applications, technologies, implementation options (e.g., physical hardware versus software virtualized function), and also architectural decisions (e.g., MEC/cache location, centralized versus distributed architecture, etc.).

This Marketing Report explains some of the use cases for applying Quality Attenuation measurement and analysis during broadband network design and lab evaluation.

Introduction

The Broadband Forum instigated an initiative called Quality Experience Delivered (or QED) [1]. This project has studied the approach known as Quality Attenuation (or ΔQ , pronounced “Delta Q”). ΔQ provides a mathematically robust performance measurement and analysis technique that integrates both packet loss and packet delay, instead of treating them separately [2]. In addition to its use in measuring and analyzing the performance of live ‘production’ broadband networks, Quality Attenuation techniques can also play a significant role in the architecture, design, and lab evaluation phases of the broadband network and services lifecycle. This document focusses on the applicability of ΔQ performance analysis to these scenarios.

1 References and Abbreviations

1.1 References

The following references are of relevance to this Marketing Report. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this Marketing Report are therefore encouraged to investigate the possibility of applying the most recent edition of the references listed below.

A list of currently valid Broadband Forum Technical Reports is published at www.broadband-forum.org.

Document	Title	Source	Year
[1] MR-452.1	<i>Motivation for Quality Verified Broadband Services</i>	BBF	2019
[2] TR-452.1	<i>Quality Attenuation Measurement Architecture & Requirements</i>	BBF	2020
[3]	<i>Analysis and predictive modelling of the performance of the Atlas TDAQ network, Lucian Leahu, PhD Thesis</i>	CERN	2013

1.2 Abbreviations

The following terminology is used throughout this Marketing Report.

ALT	Application Layer Test
BNG	Broadband Network Gateway
CDN	Content Delivery Network
CNF	Containerized Network Function
CPU	Central Processing Unit
DOCSIS	Data over Cable System Interface Specification
GPON	Gigabit Passive Optical Network
MEC	Multi-access Edge Compute
NOC	Network Operations Center
PPP	Point to Point Protocol
QTA	Quality Timeliness Agreements
ULL	Ultra-Low Latency
UX	User Experience
vBNG	Virtual Broadband Network Gateway
VDSL	Very high-speed Digital Subscriber Line
VNF	Virtual Network Function
XGS-PON	10 Gigabit Symmetric PON

2 Applicability of ΔQ to Lab Emulation

There is a pressing need to better understand the interaction of network characteristics and application user experience (UX), in order to optimize applications and conversely to audit or tune network services for specific purposes. This can be achieved by reproducing the real-time, stochastic, and instantaneous characteristics of end-to-end network transport (ΔQ) and allowing these to be varied to see their impact on application performance. This permits UX risks to be identified and managed during the design, development, and delivery of any distributed application.

Use of ΔQ performance analysis in this context provides an environment to test services, apps, etc. under controlled conditions, including reproduction of real-world measurements and credible worst-case scenarios, using the characterization of networks measured and modelled using the ΔQ framework. This supports lab-based regression testing of application performance evaluation, e.g., before deploying upgrades, and aids hypothesis-testing about the cause of undesirable behaviors and conditions.

2.1 Quality Attenuation Emulator

Quality Attenuation performance measurement (in the live network or lab) and associated ΔQ analysis can be used to characterize the performance of different broadband access technologies. The idiosyncrasies of the physical transmission layer differences between, for example, GPON, DOCSIS, VDSL, XGS-PON, and G.fast, manifest themselves as differences in IP layer performance, which is what the transported application/service experiences.

Having characterized the Quality Attenuation of a particular access technology, it is feasible to then emulate the same statistical distribution of packet loss/delay in a lab environment. Emulating the impairments experienced at the IP layer avoids the time and cost of constructing and then modifying the underlying layers and transmission systems. The “programmable” degradation (ΔQ) can then be adjusted to explore the point at which applications or services will deliver unsatisfactory QoE. The ΔQ lab capability can provide a captive experimental environment to address questions such as:

- Will my legacy network support this new application/service?
- How will my new network architecture/technology impact existing application performance?

2.2 Lab Emulation Environment

A lab-based ΔQ emulation environment should allow separate configuration of the characteristics of the upstream and downstream paths, as indicated in Figure 1 below.

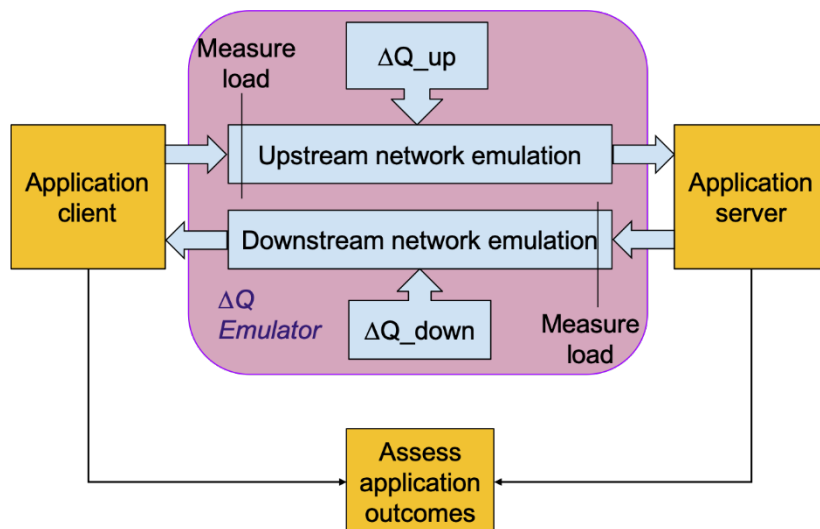


Figure 1: Lab emulation

In each direction the emulation should include:

- Rate/capacity limits
 - Including accounting for all encapsulation/framing overheads
- Geographical delay, $\Delta Q|_G$
 - Including both signal propagation and media access delays
- Serialization delay, $\Delta Q|_S$
 - Including the effect of multiple (de)serialization hops
- Variation in delay and probability¹ of loss, $\Delta Q|_V$
 - Modelling statistical effect of intra- and inter-stream contention

Each of these components can be stationary or time-varying (for example, a step-change in $\Delta Q|_{G,S}$ would represent the effect of a route change), thus providing a better laboratory environment that is more adaptable to fit desired scenarios and emulate real-world (or proposed) networks.

In addition, the ΔQ emulation environment can measure the load imposed by the application in each direction to assist in the formulation of appropriate Quality Timeliness Agreements (QTAs).

2.3 A/B testing

A further useful capability is for the ΔQ emulation environment to enable contemporaneous A/B-testing; the ability to test the impacts of different network conditions against one another in real-time (including “control” conditions), while keeping other factors such as server load constant. This is illustrated in Figure 2.

¹ While the other components of ΔQ can in principle include loss, in practice it is sufficient to model it here.

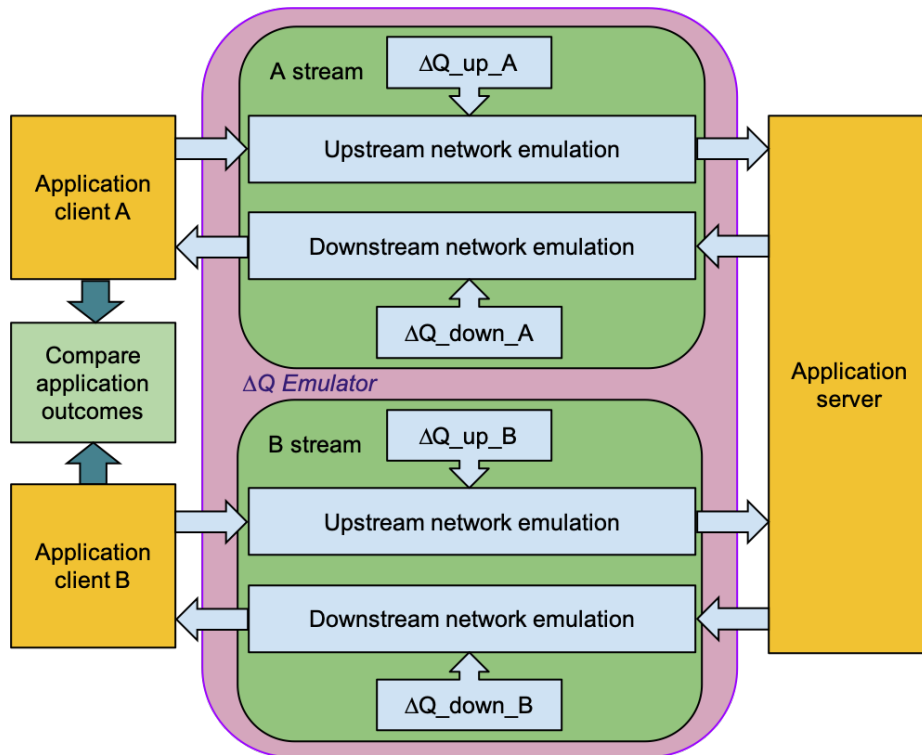


Figure 2: A/B testing

3 Δ Q Design and Lab Evaluation Use cases

3.1 De-risking Application Deployment

Application developers have various techniques at their disposal to deal with network delays and losses and to adapt to their variability. However, deciding which techniques are most appropriate for a particular application, and how much UX risk remains, can be evaluated using a Δ Q emulator:

- Using *different* variants of the application client with the *same* level of quality attenuation applied to each to see which technique works best;
- Using the *same* variant of the client with *different* levels of quality attenuation (and its variability) to find the point at which UX deteriorates;
- Reproducing the measured performance of real-world networks that the application is intended to work over, to confirm that the right UX is delivered prior to deployment.

3.2 Optimizing Speed Tests

Any network 'speed test' is simply an application, typically one in which a client transfers data to and from a server. While this is a simplistic way to measure network performance, it is widely used as a measure of capacity available to an endpoint. What are the other factors, besides the limiting capacity, that affect a speed test result? What could an ISP do to ensure that speed testers accurately report the capacity that is allocated to an end user?

Using a Δ Q emulator in its A/B configuration allows a thorough investigation of what turns a 'good' speed test result into a 'bad' one and provides guidance to ISPs and telcos of what to do (and what to avoid) to ensure they rank well in comparative speed test charts.

3.3 Measuring Cost/Benefit of Multi-access Edge Computing (MEC)

The idea of improving application performance by moving 'content' closer to the user is well-established in the form of Content Delivery Networks (CDNs) and network caches. There is increasing interest in improving performance (for example of AR/VR applications) by moving computing resources closer to the end user, i.e., closer to the 'edge' of the network. This is referred to as MEC – Multi-access Edge Compute. This proposition would benefit from having a way to demonstrate the UX benefit (in order to help justify the business case for MEC). This can be done using a Δ Q emulator, with the 'A' path simulating the quality attenuation (latency) to a central 'cloud' server, and the 'B' path the lower attenuation to the same server located closer to the network edge.

3.4 Physical Equipment Performance Benchmarking for Selection

Quality Attenuation measurements can be used to establish the relative performance characteristics of different models of network equipment which can then be used for comparison purposes. For example, a lab set up with a simple network can be used to measure the Δ Q when the network is operating with different types (models/vendors) of equipment that relay the IP packets (switch, router, Access Node, BNG etc.) [3]. The Δ Q components are derived from a number of measurement samples using a linear statistical dependency of the delay on the packet size as described in [2]. From this, the Structural Delay components

of the performance attenuation introduced by the network device (known as $\Delta Q|_G$ and $\Delta Q|_S$) can be calculated.

The serialization delay $\Delta Q|_S$ caused by an item of network equipment depends on the link bandwidth and the speed at which packets are clocked in and out of equipment (via the NIC) at network nodes. Example values of delay per byte across a single store-and-forward hop are shown below:

Link Bandwidth / Interface Speed	Delay per Byte	Difference in delay between 0 and 1500 Byte Packets
10 Mbit/s	800 ns	1.2 ms
100 Mbit/s	80 ns	0.12 ms
1 Gbit/s	8 ns	12 μ s
10 Gbit/s	0.8 ns	1.2 μ s
100 Gbit/s	0.08 ns	0.12 μ s

Table 1: Equipment interface speed and resulting delay to bits & packets

The Structural Delay of the network device encompasses its immutable properties, e.g., access to the medium, serialization/de-serialization delays, and switching speed. This approach can therefore be used to establish which equipment meets requirements (in the sense of introducing low additional latency due to packet processing). Each and every element in the packet path contributes some ΔQ even under minimal load - $\Delta Q|_{G,S}$ characterises this. The preferred solution may be one that minimises ΔQ under typical or even worst-case loads, depending on requirements. Service providers could use this measurement as part of their RFQ process to select equipment or establish its compliance against their requirements.

The use of network tomography can also be applied here. Multiple measurements can be used to deduce the performance of the network device under test which is the “unknown” entity and is simply treated as a black box that will delay IP packets. A differential paths approach can be used whereby we identify pairs of measurement points, i.e., network interfaces which communicate between them on path segments sharing the network devices. By measuring $\Delta Q|_G$ and $\Delta Q|_S$ for these paths we are able to extract the Structural Delay property for the “network device difference” (from §5.7.1.2 of reference [3]).

3.5 Virtual Network Function (VNF) Evaluation & Abstraction

Network functions are becoming increasingly implemented in software as Virtual Network Functions (VNFs) or Containerized Network Functions (CNFs). The relative performance of VNFs and CNFs from different vendors can be compared in multiple ways including throughput (where relevant), energy consumption, and resource consumption (e.g., number of CPUs). The ability to measure relative ΔQ performance gives an additional perspective for network operators seeking to select appropriate VNF/CNFs for deployment. These software-implemented functions could be evaluated individually (such as authentication, PPP termination, routing) or in combination (e.g., as the main elements in the packet flow data path of a BNG user plane).

The measurement of ΔQ enables us to reason about the performance characteristics of a network in a way that abstracts us away from the implementation details of the network. This abstraction can be at various levels depending on our chosen granularity for the observation (measurement) points. It can for example be applied to international, national, local, or home networks. It can also be applied to individual network links and elements, such as switches and routers or a combination of VNFs/CNFs. The $\Delta Q|_G$ and $\Delta Q|_S$ components could be used to compare the performance of various VNF/CNF types, vendor implementations, or to compare physical versus virtualized implementations. The $\Delta Q|_V$ component can provide useful information to understand scheduler properties.

An example of abstracting ΔQ to a higher level (such as a network service consisting of a chain of VNFs/CNFs) is illustrated below. Here, a system such as a network node involves use of network,

processing, and storage resources. This could be typical of a virtualized network node implemented on generic compute hardware. It has three main tasks: Receive some data from the network, perform some compute processing on that data, and then store the outcome. We can measure ΔQ outcomes for each of these constituent tasks or we can choose to abstract the ΔQ to a single global outcome treating the three resource domains in the network element as a single aggregate resource. Note that, because ΔQ is conserved, if the aggregate ΔQ is within some designated limit, then all the ΔQ s for the constituent tasks must also be suitably bounded.

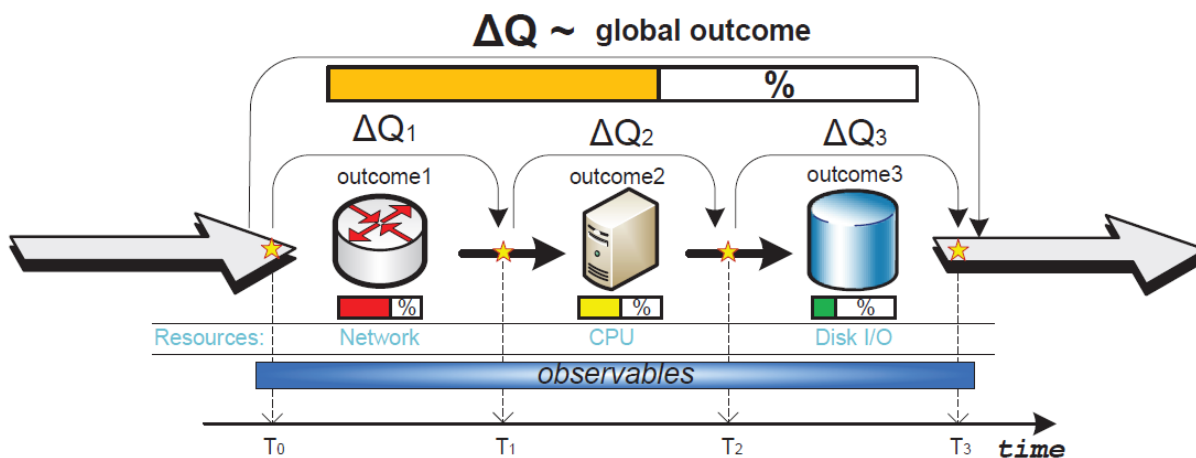


Figure 3: Example of abstraction of ΔQ observable outcomes

Measuring performance and resource consumption for every function, node, and link, and a large network could become complex and unwieldy. Hence abstraction can be extremely useful if it enables the understanding of performance outcomes pertinent to the use-case in question without incurring undue complexity and cost. Likewise, modelling all the resource access and contention mechanisms in the large shared resource environment of a broadband packet network can become unmanageable. For scrutinizing the performance of a particular task or network element (virtualized or physical) it can be useful to abstract away from the other elements and processes in the network.

3.6 Equipment Configuration Impact

The preceding use-cases have shown that ΔQ can be used to benchmark the performance of different physical or virtual network functions, e.g., from different vendors as part of a commercial tender. It can also be used to evaluate the performance impact of different configurations of equipment/VNFs. Two examples are described below for GPON access transmission equipment.

3.6.1 Comparing GPON Configuration Impact

ΔQ can be used to analyze the performance impact of different GPON configurations, QoS profile parameters and performance variation under load (e.g., Dynamic Bandwidth Allocation). The results in Figure 4 below compare latency measures of two different OLT configurations.




Figure 4: Comparing GPON Configuration Impact
 (Green= $\Delta Q|_G$, Blue= $\Delta Q|_S$, Red= $\Delta Q|_V$ Mean, Amber= $\Delta Q|_V$ Standard Deviation)

In the chart on the left, the GPON scheduler is acting to keep the average round trip latency constant; the packet serialization ($\Delta Q|_S$) component in blue is constant; the $\Delta Q|_G$ component (green, transmission latency) is changing but the average $\Delta Q|_V$ component (scheduling, queueing, buffering latency) is adapting to try and make the overall average latency (top of the red line) constant.

In the chart on the right, the overall average latency results are also quite flat but not as “tight”. The OLT operates in L2 forwarding mode (not based on TCP connection). The upstream and downstream have independent forwarding; the current forwarding packet and the next forwarding packet are independent, hence the OLT can’t adjust the return delay to a constant purposefully.

OLT configuration differences can result in a difference in standard deviation of $\Delta Q|_V$.

For the OLT generating measurement results on the left:

- Different ONTs use different DBA, different group profiles, even under one GPON port many ONTs use different configurations.

For the OLT generating measurement results on the right:

- All the VOIP services use fixed DBA (type1 fixed 1024); and for the internet service (type3 assured 1024, max 10240)

3.6.2 Impact of Turning on QoS on GPON

Isolating the $\Delta Q|_V$ component (by post-processing measurement results), it is possible to focus on the impact of scheduling on performance.

In this experiment on a GPON line, the connection was streaming video traffic in both directions before QoS was enabled. By plotting the separate upstream and downstream directions of $\Delta Q|_V$, Figure 5 shows the improvements achieved after the activation of QoS on GPON technology.

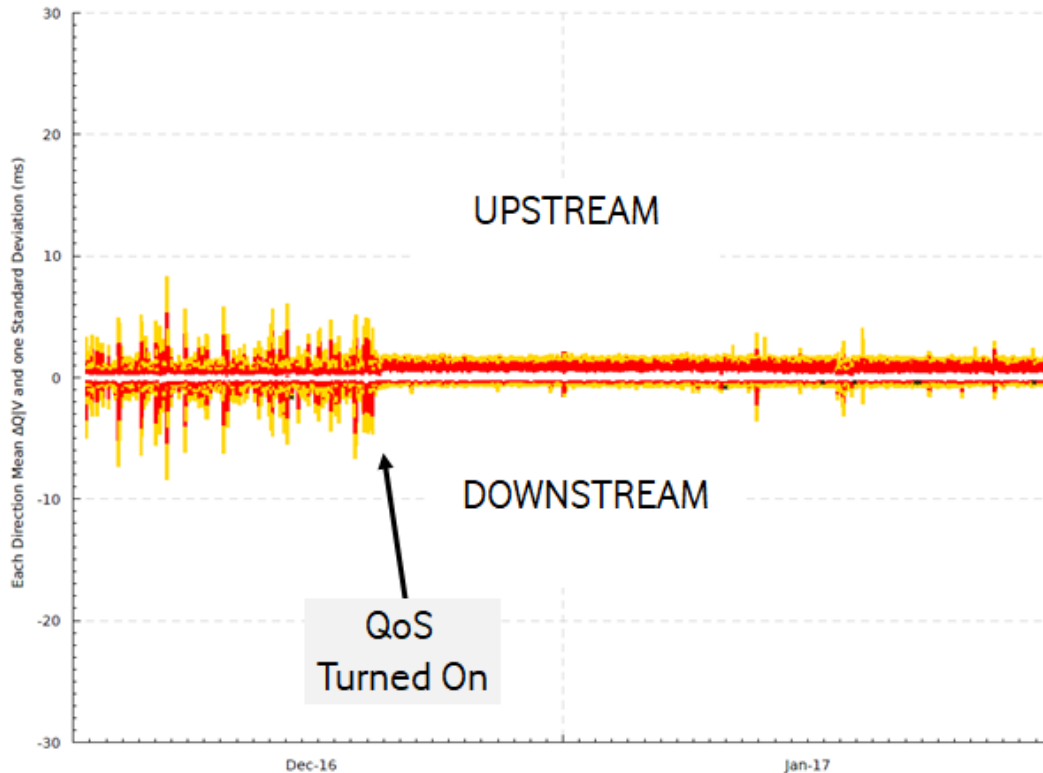


Figure 5: Impact of turning on QoS on GPON
 (Red= $\Delta Q|v$ Mean, Amber= $\Delta Q|v$ Standard Deviation)

3.7 Input to Network Architecture Design/Analysis

Quality Attenuation performance analysis can be used for measurement in live networks, measurement in the lab, and also in mathematical models. The latter can be used to reason about the performance implications of different network architectures and design approaches before the network is built and can help to give answers before costly trials are deployed. In turn, lab and field measurements can be used to inform the models to refine them and to undertake “what if” scenario analysis. For example, a network model informed by ΔQ measurements is able to provide a predictive capability in case different operating conditions are experienced, e.g., to evaluate the impact of increased network traffic loading and to quantify the consequent increased delays in the network. Hence this can provide a useful capability for network architects and designers.

The architecting and design (including dimensioning) of networks often relies on average values like bandwidth. However, the required loss and delay performance metrics (which are instantaneous characteristics previously difficult to predict at design stage) do have an impact on the scalability of the network and also on its behavior close to saturation. For network architects, Quality Attenuation can provide an end-to-end network performance description, using its piece-wise convolution properties for the performance contribution of the various links and nodes (network devices) that the network is comprised of. It can also provide a measure of the delay caused by queueing. Using this information in a comparison with theoretical traffic queueing models can enable the performance to be inferred for different traffic loads and traffic arrival patterns.

Some examples of the sort of architecture and design decisions that can be informed by Quality Attenuation modelling and measurements include:

- Performance implications of locating a Virtual Network Function (VNF) located at different nodes within the network or Internet. For example, on virtual CPE versus the network edge (as per MEC) versus a centralised private cloud (network operators data centre) versus a public cloud.
- As above but for a content cache or CDN node.
- Performance of different network slices, e.g., “standard” versus Ultra-Low Latency (ULL).
- The ability to compare the relative performance of different network equipment implementation approaches. For example, comparing an ASIC-based element (such as physical BNG data plane) versus an x86-based data plane used by a vBNG VNF.

3.8 Use of ΔQ in Systems Integration

When accustomed to large-scale systems that work well, such as the telephone system (PSTN) and mobile networks, it is easy to overlook how much effort those systems required to make their operation robust, i.e., to move from a service that ‘mostly works’ to one that rarely fails - a system on which users can rely. The difficulty is that the interaction of even a small number of relatively simple components can result in emergent behavior that is unexpected and difficult to predict or control. This is particularly the case where matters of performance and timing are concerned; purely functional issues are straightforward by comparison. Systems that exhibit the desired success modes in a development setting will often manifest unexpected failure modes when deployed in the field; this is largely due to differences in performance of various components in real world deployments. For example, changes in the response time of one component of the system can negatively impact the performance of others in ways that were not exposed in the development setting.

Consider, for example, three different implementations of a 3GPP network: the situation of macrocells backhauled over TDM circuits; IP connectivity (for microcells in a geographically small area) over Ethernet; and femtocells backhauled using ubiquitous and cost-effective IP over broadband, such as xDSL. The key point is that these three different systems have the same *success modes*, i.e., when they work they do so in the same way, according to the 3GPP standards; however, they are likely to have *different failure modes*, due to the different performance and timing properties of system components.

Some of the hazards² in the system will depend on its scale, both in space and in time. As the extent of the system grows, new behavior may emerge, for example because different components of the system scale differently and interact differently with each other. On the other hand, new hazards may emerge as time passes, because of changes in the behavior of system components or in the operating environment.

There are two ways of tackling this problem:

1. Starting with each component, determine the performance it requires from the rest of the system, and the performance it delivers provided those requirements are met. Then iterate between all the components to see whether or not they can collectively meet each other’s performance needs (this will depend on a range of external factors such as load). This approach is easiest to apply when there is a clear layering of the system and the requirements and service of each layer on those above and below it are easy to quantify;
2. Starting with the complete system, measure its performance in a range of conditions so that the arming of hazards can be observed. Then ‘divide and conquer’ to discover which interaction between components is arming the hazard. In this way emergent hazards and their arming conditions can be discovered.

² A ‘hazard’ is the potential for a negative or harmful outcome; it is considered ‘dormant’ if no negative consequences can occur under present circumstances; ‘armed’ if they can occur (i.e., there is a risk); and ‘active’ if they are actually occurring. For example, since packet buffers are always finite, the hazard of packet drops always exists, but this is only *armed* by the presence of packets (empty buffers can’t drop anything), and only *active* if drops occur.

Characterizing the ΔQ properties of individual components gives a lower bound on the ΔQ we can expect when they are combined together. Comparing this with the desired ΔQ can save considerable Systems Integration effort by identifying infeasible solutions/configurations in advance. Likewise, comparing the ΔQ delivered by a subsystem with that expected from the combination of the ΔQ s of the components can expose errors in the subsystem construction/configuration.

4 Summary and Conclusions

Quality Attenuation for performance measurement and analysis can be used to assess “in-life” broadband network performance but can also play an extremely valuable role in other phases of the technology lifecycle as illustrated below. These include design and lab evaluation of new equipment, applications, technologies, implementation options (e.g., physical hardware versus software virtualized function), and architectural decisions (e.g., MEC/cache location, centralized versus distributed architecture, etc.).

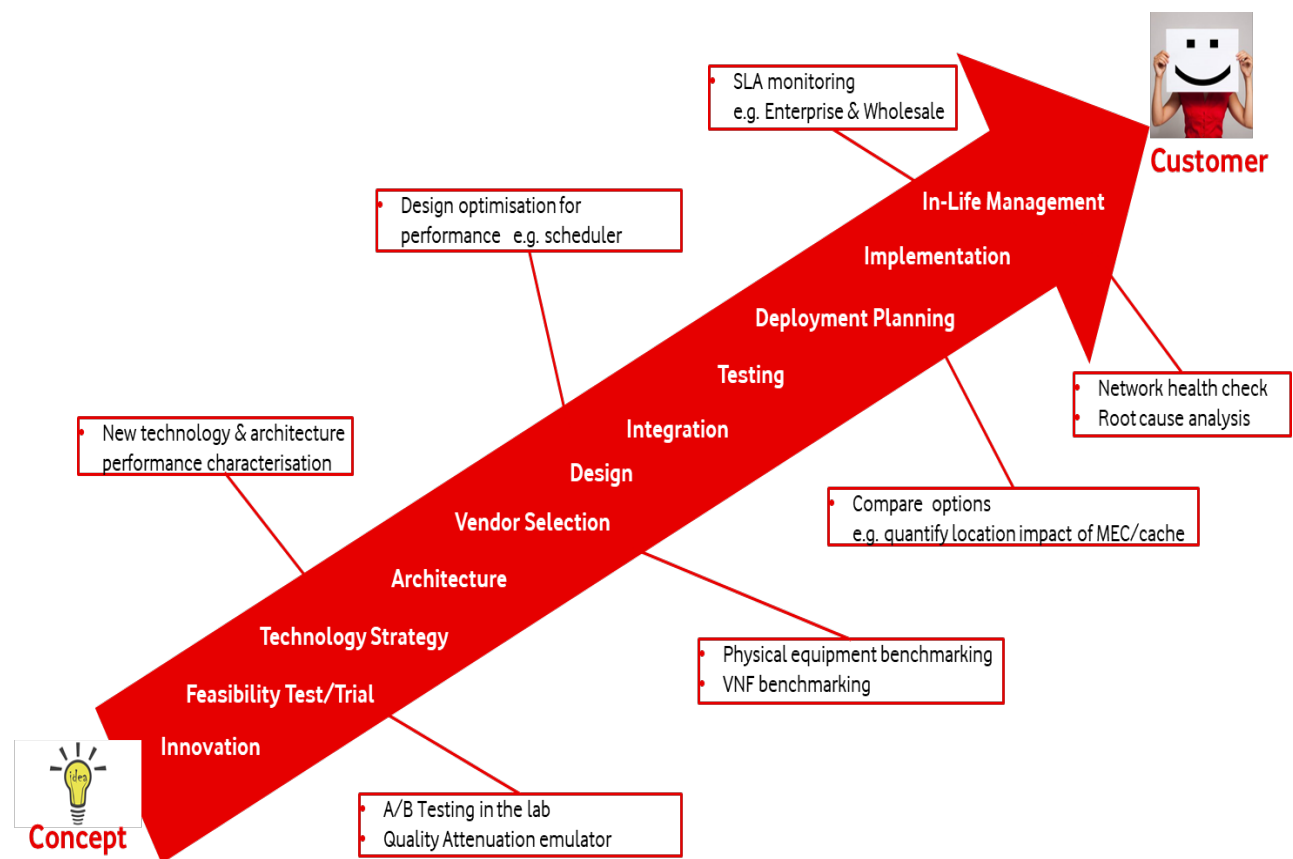


Figure 6: Potential roles of Quality Attenuation measurements in the technology lifecycle

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Marketing Report has been approved by members of the Forum. This Marketing Report is subject to change. This Marketing Report is owned and copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Marketing Report may be owned and/or copyrighted by Broadband Forum members.

Intellectual Property

Recipients of this Marketing Report are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of this Marketing Report, or use of any software code normatively referenced in this Marketing Report, and to provide supporting documentation.

Terms of Use

1. License

Broadband Forum hereby grants you the right, without charge, on a perpetual, non-exclusive and worldwide basis, to utilize the Marketing Report for the purpose of developing, making, having made, using, marketing, importing, offering to sell or license, and selling or licensing, and to otherwise distribute, products complying with the Marketing Report, in all cases **subject to the conditions set forth in this notice and any relevant patent and other intellectual property rights of third parties (which may include members of Broadband Forum)**. This license grant does not include the right to sublicense, modify or create derivative works based upon the Marketing Report except to the extent this Marketing Report includes text implementable in computer code, in which case your right under this License to create and modify derivative works is limited to modifying and creating derivative works of such code. For the avoidance of doubt, except as qualified by the preceding sentence, products implementing this Marketing Report are not deemed to be derivative works of the Marketing Report.

2. NO WARRANTIES

THIS MARKETING REPORT IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT AND ANY IMPLIED WARRANTIES ARE EXPRESSLY DISCLAIMED. ANY USE OF THIS MARKETING REPORT SHALL BE MADE ENTIRELY AT THE USER'S OR IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY USER, IMPLEMENTER, OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS MARKETING REPORT, INCLUDING BUT NOT LIMITED TO, ANY CONSEQUENTIAL, SPECIAL, PUNITIVE, INCIDENTAL, AND INDIRECT DAMAGES.

3. THIRD PARTY RIGHTS

Without limiting the generality of Section 2 above, BROADBAND FORUM ASSUMES NO RESPONSIBILITY TO COMPILE, CONFIRM, UPDATE OR MAKE PUBLIC ANY THIRD PARTY ASSERTIONS OF PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS THAT MIGHT NOW OR IN THE FUTURE BE INFRINGED BY AN IMPLEMENTATION OF THE MARKETING REPORT IN ITS CURRENT, OR IN ANY FUTURE FORM. IF ANY SUCH RIGHTS ARE DESCRIBED ON THE MARKETING REPORT, BROADBAND FORUM TAKES NO POSITION AS TO THE VALIDITY OR INVALIDITY OF SUCH ASSERTIONS, OR THAT ALL SUCH ASSERTIONS THAT HAVE OR MAY BE MADE ARE SO LISTED.

All copies of this Marketing Report (or any portion hereof) must include the notices, legends, and other provisions set forth on this page.

End of Broadband Forum Marketing Report MR-452.4