PREDICTABLE
NETWORK
SOLUTIONS

# *RESPONSE TO "DESIGNING THE BROADBAND UNIVERSAL SERVICE OBLIGATION"*

23rd June 2016

Predictable Network Solutions

We see the proposal for a USO as a pivotal moment in the evolution of our increasingly digital society. It creates an opportunity to specify a service that will be of lasting value to citizens, sidestepping the need to compete on (frequently irrelevant) service features.

This contribution addresses the question "How should the minimum technical performance of the USO be specified?", in particular the parameters in which this specification should be couched. This is important because the utility of a service depends on more than the 'broadband speed', which is simply a measure of the average rate at which bits are transmitted; even a simple 64kbits/s PSTN service is not fully characterised without considering availability and call blocking probability. It is not uncommon "high speed" services to fail to deliver an acceptable experience - it is not the "speed" that is lacking, but "the quality".

For a statistically multiplexed broadband service there are many ways in which the quality of the service can vary independent of the bit-rate. For example, the figure below shows measurements of individual end-to-end packet delays (sorted by packet size) over two ADSL lines, both of which synced to exactly the same rate. The frequency of large delay outliers across Connection B makes it far less effective at delivering a VoIP service than Connection A.

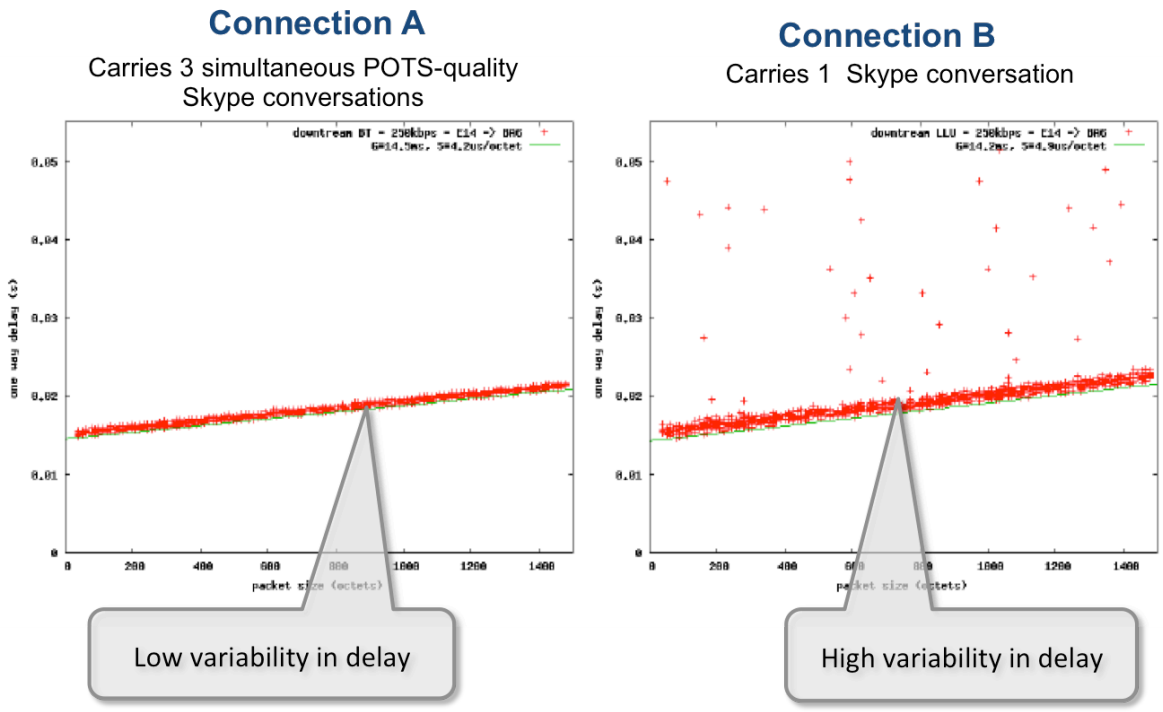


Transporting packets with low loss and low delay consumes resources in the network, however, and thus has a cost. Thus:

1. There is a limit to the quantity of packets (the offered load) that can be transported with a given level of loss and delay;
2. Specifying the loss and delay too tightly will make the service very costly to deliver.

Our proposal is to specify the USO as a '***quality floor***', i.e. a *minimum level of quality of packet transport* that will be maintained provided the *offered load* does not *exceed* a specified level. Note that the bound on the offered load may differ from the maximum rate that the physical connection supports; how load in excess of the USO rate is treated need not be specified by the USO. Note also that some network resources are consumed on a per-packet basis, independently of the size of the packet, and so it may be important to place a bound on the rate of packets as well as the aggregate bit-rate of those packets. The quality floor would create a new differentiator in the market; one that can be used uniformly across all delivery technologies and provide a product differentiation that has a direct correlation to end-user benefit.

Another key point about the delivered quality of a statistically-multiplexed packet service is that it only gets worse, in the sense that if a packet experiences delay it cannot be 'un-delayed' and if it should be lost it cannot be 'un-lost'. 'Quality' is thus a negative attribute like darkness or silence, and so we need to consider its converse that we call 'quality attenuation', which is roughly like noise in an electrical circuit. Delivering a quality floor is thus a matter of ensuring that quality attenuation is bounded.

This proposed quality floor begs the question of how quality attenuation is to be specified and measured. The essential point here is that average measures are insufficient. Consider the figure above: the average packet latency for Connection A and Connection B is almost exactly the same, and yet Connection B has far lower utility. Note that 'jitter' is not a useful measure either, since it refers only to the difference in delay of successive packets (consider an example where every packet has the same delay, except that one packet in twenty has 10x the delay of the others: the average jitter will be 5%, which seems quite low, but this will not be a good service for many applications). An appropriate measure does, in fact, exist (called ΔQ) and is described in Appendix A of the 2015 Ofcom report "Traffic Management Detection Methods & Tools" (http://stakeholders.ofcom.org.uk/market-data-research/other/technology-research/2015-reports/traffic-management).

The delivery of a quality floor must be specified and measured between particular reference points, since the service to an arbitrary point in the Internet depends on many factors that are outside any USP's control. One reference point will be the service handover point from the USP to the end-user premise; we suggest (for the purposes of defining the USO) the other could be the nearest UK national capital.

Since quality attenuation accumulates over the end-to-end path, and it is the end-to-end attenuation that determines the end-user experience, the USO quality floor must be specified so that reasonably expected levels of quality attenuation in other parts of the end-to-end path (premise WiFi, Internet transit etc.) can be accommodated while still delivering an acceptable overall service.

Packet delays can be divided into a constant component, which depends on factors such as distance and the speed of light, and a variable component, independent of geography and line rate. We suggest the constant component should be no more than a constant factor

times the time for light to travel between the reference points; a reasonable factor would be 2. The variable component of delay (and the loss rate) should be uniformly bounded for all end-user premises; note that this component typically has a larger impact on the delivered user experience for many applications.

The overall quality attenuation for offered loads within the USO limit should be sufficient for a range of applications important for modern digital citizenship to function acceptably well. Since the delivery is to a premise, some level of concurrent usage should be supported.

Note that if the quality attenuation bound is exceeded for applied loads below the USO load limit than the USO specification has been broken. Thus a basic level of monitoring that the USO is being delivered can be performed at low loads (which has a corresponding lower impact and operational cost than existing speed test measurements).

Over time, it may become desirable to raise the load limit for the USO to accommodate new essential applications, but this should not increase the quality impairment, otherwise previously functioning applications may start to fail. The quality impairment floor explicit in such USO definition would also have beneficial on the economic ecosystem; app developers would have a known worst case to design and build their innovative products against.

It is likely to be uneconomic to deliver the specified bounds on quality attenuation at the maximum allowed load to all endpoints at all times. Thus a level of compliance can be set for a USP, but this should be on a *per-endpoint basis* (current measurement approaches are over an average of all end points). Thus a 90% compliance would mean that all endpoints experienced the specified service 90% of the time, not that only 90% of the served endpoints received the service (which would be a failure of universality).