

# Overcoming the challenges posed by using indirect inference when solving performance problems on the Internet

October 10, 2017

By Frank Mungau, Ph.D.

Whereas congestion is readily apparent to *Internet subscribers* when their download speeds deteriorate resulting in performance degradation such as a lower video quality and/or buffering, the same cannot be said of *Network Operators* who have to find a way to detect congestion based purely on observations of the flow of data traffic across the network.

In terms of the links between network nodes (e.g., routers), the network can be said to be congested if the traffic demanded by subscribers (measured in Mbps for example) is arriving at a rate that is greater than the transmission rate of the link which has a fixed speed or “capacity” (also measured in Mbps). However, the dominant data Transport Protocol on the Internet, the so-called Transmission Control Protocol (TCP) confounds such a straight forward definition because TCP is designed to maximise the rate at which information is sent whilst constantly probing for the maximum rate that *triggers congestion*.

What usually happens when the data demand exceeds the link capacity (causing congestion) is that the excess traffic is stored in buffers where it is delayed until the link has cleared the backlog; or in extreme cases when the buffer is already full, the excess traffic is not saved (i.e., it is dropped) and is thus lost in transit. In the first instance, TCP is designed to rapidly increase the rate at which it sends individual data packets during a bulk data transfer (the *slow-start* mechanism) until it receives an implicit control message such as “no” acknowledgement from a receiver indicating a dropped packet, or a delayed acknowledgement indicating increasing delays. Once this occurs, TCP takes measures to retransmit any lost packets and will eventually enter *congestion avoidance* when it reaches some fraction of the sending rate at which congestion occurred, thus stabilising the delays and packets losses across the congested link. In the case where there is more than one subscriber sharing the link, the flow of traffic associated with a given user will try to push other competing flows out of the way to claim its share of the available capacity when TCP attempts to fill the buffer with a given user’s payload (contention) until (hopefully) other flows experiences drops or delays (congestion).

Because data is transmitted using TCP as described above, occasional congestion episodes of moderate duration and frequency are *inevitable*. The *duration of the congestion episode* is therefore an important factor when evaluating its *potential impact on users*. But in order to solve performance problems, it is necessary to quantify congestion which can be done based on several different definitions [[Link to MIT definitions paper](#)].

The **text book definition** relies on the consequences of congestion; i.e., a queue of packets forms at the input of a congested link causing packets to be dropped when the buffer is full (or is filling up in the case of [Active Queue Management](#), AQM). In this case, congestion is measured (for example) as the rate of dropped packets.


The **network operator definition** of congestion is the link *utilization* (a.k.a. *traffic intensity*) [[Link to University of Twente thesis](#)] which is calculated as the ratio of the amount of data traffic demanded to the link capacity. For example, a network operator might conclude that a link is congested if the observed measurements of the data demand reach 70 – 80% of the link capacity, in which case a link capacity upgrade is planned. It is important to note that the utilization as described here is necessarily defined in terms of the *average data demand* (more on this shortly) because TCP will attempt to maximise the send rate and invariably exceed the link capacity (at least momentarily) during normal operation. Another alternative definition of congestion is the **economic definition** which defines congestion in terms of the costs (not necessarily economic) that are imposed on the individual users sharing a link due to the lack of adequate resources (capacity). The economic definition is interesting because it implies sharing among different entities so that the cost is shifted from one user to another. Indeed, new directions in research [[Link to UC Berkley/CAIDA paper](#)] have investigated the statistics of packet delays for individual TCP data flows when localising congestion.

Because “some moderate” congestion is the normal state of affairs on the internet, it is clear that the ability to solve performance problems based on the definitions given above poses several significant challenges – all of which can be explored in terms of the network operator’s definition of congestion, without loss of generality:

1. First, there is the question of measuring the “average data demand” which is difficult to define because the flow of data arising from TCP’s *slow-start* and *congestion avoidance* mechanisms is anything but a steady stream of bits. In practice, what this means is that the figure for the average demand is different if calculated over a 5 minute or a 1 second interval because the total amount of traffic divided by the measurement interval yields significantly different results (see the figure below). The question is, which time period should one select when quantifying traffic intensity (...it is also true in general, that the higher the time resolution, the higher the cost of the measurement equipment)?

- Next there is the question of setting a bar (for example 70 – 80%) when determining if a link is indeed critically overloaded. In this case, it is not clear whether the percentage chosen to define an overloaded link has a direct bearing on the poor performance experienced by the user; i.e., how much congestion is too much congestion for the user to bear?
- Finally, there is the question of the *duration* that is used to determine if the congestion event is indeed significant and should warrant a link upgrade. Specifically, it is not desirable for the network operator to evaluate congestion over several hours per day, over many weeks or months, because this is too long a period over which internet subscribers may be negatively impacted.

From the arguments listed, it can be concluded that the methods described above for defining and managing congestion all amount to *indirect inference*; in other words, it is difficult to determine how each particular user is affected by observing the sum total of data traffic traversing the network.



This rather stark conclusion (that it is difficult to determine the extent to which the user experience is harmed by measuring network congestion) is supported by unambiguous statements by various prominent authorities including David Clark from MIT [\[Link to MIT/CAIDA paper\]](#) when describing the *state of the art* method used in their research [\[Link to UC San Diego Blog\]](#) which was funded by the National Science Foundation (Oct 1 2014 – Sep 30 2017) [\[Link to CAIDA project page\]](#) and used to provide regulatory insight for the Federal Communications Commission.

**“Our measurement scheme can detect persistent congestion, but does not permit us to diagnose its cause, determine the extent to which it is harming the user experience, or identify how to best alleviate it.”**

Furthermore, Princeton’s Nick Feamster wrote an interesting article in the Centre for Information Technology *Freedom to Tinker* [\[Link to the Blog\]](#) which discusses how direct measurements of *aggregated* (out-of-band) delays in various networks do not adequately disambiguate the location of congestion in the Internet. Incidentally, this work is cited by authors at Princeton/CAIDA, who have devised an alternative method recently published in the literature [\[Link to Princeton/CAIDA paper\]](#) on using the statistics of the packet delays from individual flows (In-band techniques, the so-called *TCP Congestion Signatures*) to locate congestion. Despite these recent advances, the indirect inference methods remain unsuitable for determining how network congestion *quantitatively* affects the user’s experience.

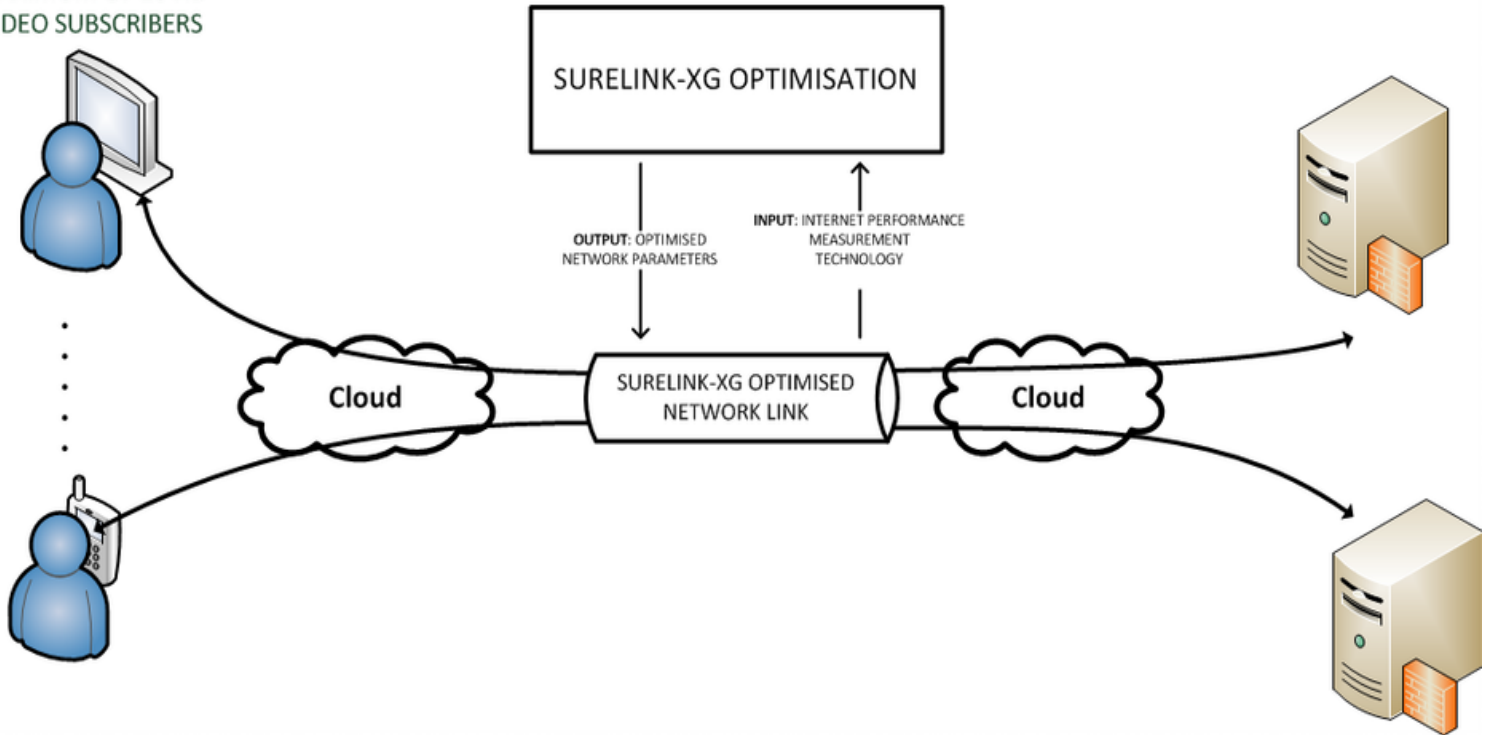
An alternative (and more *direct*) approach is to choose a measure of congestion that quantifies the user experience such as the *download/connection speed* and then determine how much network capacity is required to meet the demand. There are other measures that relate to the user’s Quality of Service on the Internet such as the time it takes to send data and receive an acknowledgement (round trip time is the two-way latency), the variation in inter-arrival time between successive packets (jitter) and packet loss due to transmission errors, congestion of links, faults or the routing process. However, connection speed is arguably *the* definitive measure of the Quality of Service because latency has been minimised in modern networks and the other two factors, packet loss and jitter are compensated for by advanced error correction codes and buffering at the receiver, respectively. Fortunately, it is possible to accurately formulate the speed of an internet connection based on variables such as the two-way latency and the packet loss ratio.

The analytical methods that exist to determine connection speed when multiple users share a link essentially depend on the way in which different TCP data flows contend for the available capacity and whether they experience packet loss due to congestion concurrently (typically the case for homogeneous flows) or at different times (heterogeneous flows) [\[Link to the Book on Internet Congestion by Subir Varma, Chapter 2\]](#). In either case, it is possible to calculate the average rate of an Internet flow from the Capacity of a link but; until recently,

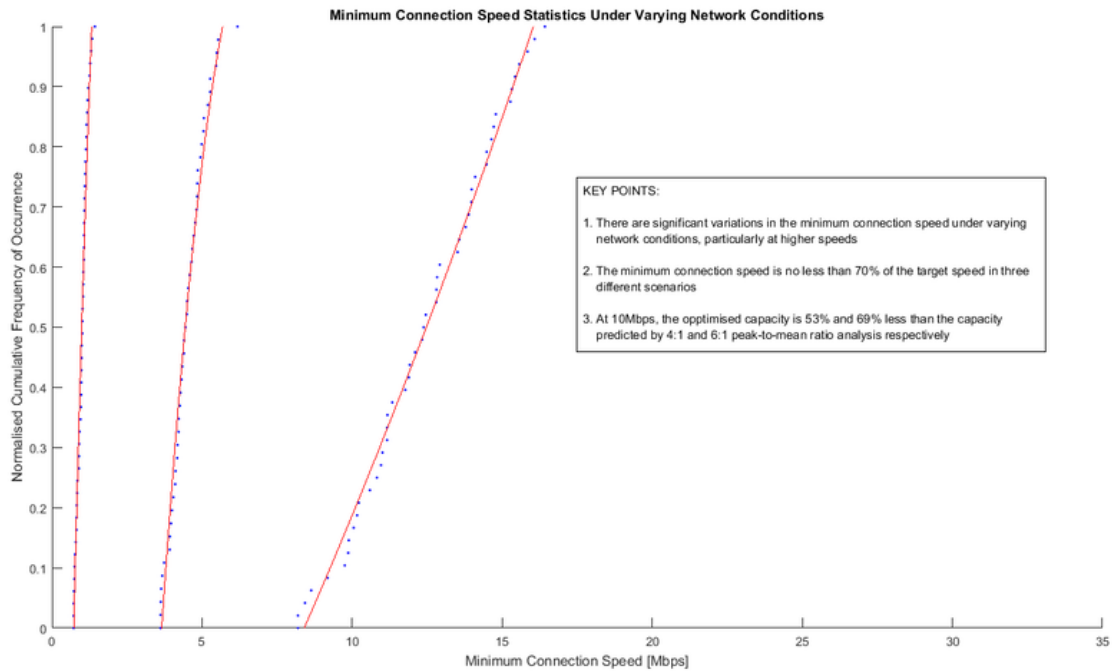
significant challenges have remained when solving this problem practically owing to the lack of a closed form expression of the relevant equations.

MAXIMUM OF 10 HD VIDEO SUBSCRIBERS

HD VIDEO SERVERS



The author of this blog has recently devised a computationally efficient method of solving the analytical equations with only a minor penalty in accuracy, which is to the best of the author's knowledge, the first (and only) record of this result having been achieved (first achieved 12th May 2017). The figures above & below depict the scenario & results respectively of a proof of concept in which the required variables are derived from network measurements and the link capacity required for users to download data at a target speed is determined. An extensive network simulation campaign shows that the measured download rates are at worst 70% of the desired rate (*speed tiers* of 1, 5 and 10 Mbps) under widely varying network conditions, the latter of which can be associated with variations in network congestion at other locations in the Internet (the network clouds). Interestingly, the utilisation of the link is well above 90% when the target speeds are universally realised.



What this important result implies is that it is now possible to optimise the Internet as a whole by engineering each and every link to ensure that it is *capable* of supporting download/upload speeds at a desired data rate. This new capability is a fundamental (first principles) step towards closing the gap between *advertised* and *measured* speeds on the Internet which will lead to greater user expectation (and thus greater user demand) of Internet services. The ability to accurately manage network congestion could also avoid unedifying scenarios which have resulted during past congestion episodes [\[Link to the Princeton Blog\]](#); incidents which are only likely to increase as the demand for high-bandwidth content surges and the Internet is further re-engineered to cater for the subsequent growth in traffic. Finally, because it is widely recognised that broadband Internet services are a key facilitator for social and economic development, the new capability will allow Internet Service Providers to abide by the regulations stipulated by national authorities such as the FCC, Ofcom and CA Kenya [\[Link to Measuring Broadband America\]](#), [Link to BBC News on Ofcom Announcement](#), [Link to Daily Nation CA KEN Announcement](#)].