Statistical Oversubscription

by Oleg Berzin ()

In the world of network capacity planning statistical considerations play a key role. It is intuitively clear that the "worst case" provisioning results in wasted network capacity and unnecessary expenses. It is also intuitively clear that in an oversubscribed network there is a risk of not being able to provide adequate services to users if the oversubscription ratio is too aggressive.

Network capacity planning is a demanding area which is oftentimes based on empirical observations of network usage trends, patterns and forecasts. In a "bursty" traffic world it is difficult to derive precise link sizing parameters and we often find networks designed for a perceived worst case usage scenario. However, it is precisely the "bursty" nature of traffic that allows using a statistical approach in capacity planning. In other words we would like to be able to exploit the fact that "not everyone is on at the same time" to make our network sizing decisions. Moreover, we would like to be able to take the "guesswork" out of our technical decisions (as much as possible) and use some fundamental statistical principals in providing answers to network planning problems. Consider the following scenario.

A number n of individual "on/off" traffic sources will be using a network communication link. Each source produces a traffic burst of up to rate R (bits per second) in a time interval of T seconds (the "on" period), then a source is "off" for another T seconds and after that time the cycle repeats. Note that a source sends RT bits during an "on" period and the number of bits per period is random. The individual sources are not synchronized in time (they can go "on" and "off" independently from each other), and their aggregate traffic is routed via the network link. We need to provide an estimate of a required link capacity considering a statistical nature of the offered aggregate traffic. The traffic source behavior is shown in the figure below.

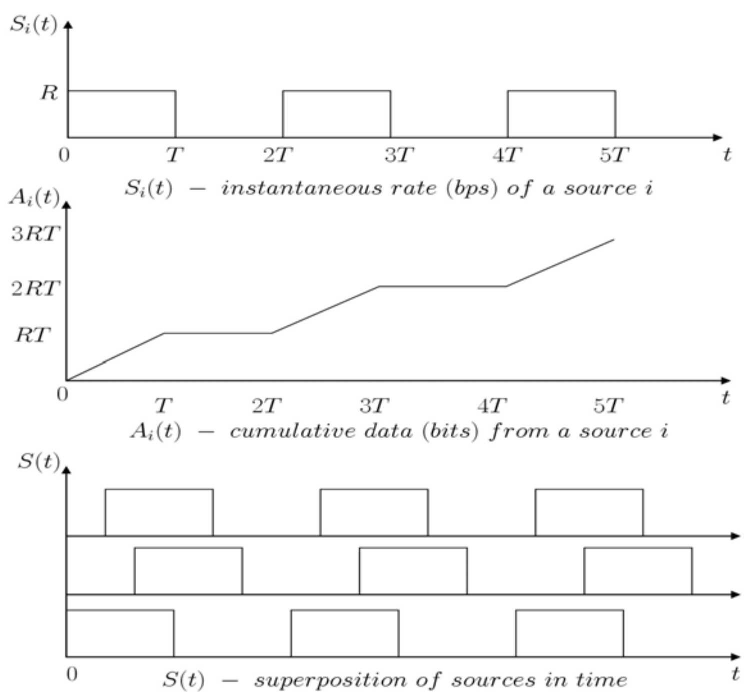

Intuitively we would expect that over time approximately ½ of all sources will be "on" at any given instant in time. From the figure above the average rate of any given source is $\frac{R}{2}$. The cumulative average rate from all sources is $\frac{nR}{2}$. In addition a rate variance (deviation from average) of a given source is $\left(\frac{R}{2\sqrt{3}}\right)^2$, and an aggregate rate variance is $n\left(\frac{R}{2\sqrt{3}}\right)^2$. The rate standard deviation of an individual source $S_{max} = \frac{R}{2\sqrt{3}}$ captures the variability of the rate with which a source is sending its data during the "on" period. The computation of S_{max} above is based on the assumption that a given source may vary its rate over time interval T in a uniform manner (with equal probability during the transmission interval).

The statistical analysis for this type of problems is usually based on the fact that a large number of independent traffic sources over time will tend to behave according to the Gaussian (bell curve) distribution with mean rate $\mu = \frac{R}{2}n$ and rate standard deviation $\sigma = \sqrt{n}S_{max}$. This is based on a statistical law often referred to as a Central Limit Theorem (CLT).

To apply the CLT to our problem, consider a Normally distributed random variable Z that represents the aggregate traffic rate from all n sources. Also consider a value C that represents the network link capacity (bps) for the link to be used by the traffic from our sources. Since according to CLT the aggregate rate is a Gaussian random variable ($\mu = \frac{R}{2}n$ and $\sigma = \sqrt{n}S_{max}$) we can normalize it and express it in terms of Z - a Normally distributed random variable ($\mu = 0$ and $\sigma = 1$):

$$Pr\left(Z > \frac{C-\frac{R}{2}n}{\sqrt{n}S_{max}}\right) = 1 - \Phi\left(\frac{C-\frac{R}{2}n}{\sqrt{n}S_{max}}\right)$$

Expression above gives a probability of an aggregate source rate exceeding link capacity C where the aggregate rate is represented by a Gaussian random variable with mean and standard deviation of $\mu = \frac{R}{2}n$ and $\sigma = \sqrt{n}S_{max}$ respectively. This expression makes use of the normalized Gaussian random variable Z and the Normal cumulative distribution function $\Phi(x)$. To make use of the statement above consider the following requirement: we would like to ensure that the probability of an event when the aggregate source rate exceeds link capacity is acceptably low and is equal to a small ϵ. Or in the mathematical notation:

$$1 - \Phi(C_\epsilon) = \epsilon$$

Where, $C_\epsilon = \frac{C-\frac{R}{2}n}{\sqrt{n}S_{max}}$ may be referred to as a quality of service grade. The QoS Grade reflects the confidence interval with which we can declare that the resulting capacity will not be exceeded by the offered traffic load. Given the relationship above we can write:

$$C = \frac{R}{2}n + C_\epsilon S_{max}\sqrt{n}$$

The equation above is the Statistical Oversubscription formula that says that the planned link capacity C should be based on the aggregate average rate ($\frac{R}{2}n$) plus a statistical over-subscription factor (SOSF). The SOSF accounts for a rate variability of each source (S_{max}), the number of sources ($\sqrt{n}$) and a QoS Grade (C_ϵ). Note that SOFS is proportional to the square root of the number of sources and that the link capacity computed using this formula yields a value that grows slower with n than the worst case capacity of nR.

To illustrate the use of this formula consider the following example:

$$n = 100,\ R = 1{,}000{,}000\ bps,\ T = 1\ sec, \epsilon = 0.01$$

For the worst case provisioning scenario we require that:

$$C_{worst\_case} = nR = 100 \times 1{,}000{,}000\ bps = 100\ Mbps$$

For statistical over-subscription, first let's compute the QoS Grade using the Normal Distribution Table (http://en.wikipedia.org/wiki/Normal_distribution). The value of $\epsilon = 0.01$ corresponds to a confidence interval of 99% $(1 - \epsilon)$, which in turn corresponds to a value of the QoS Grade $C_\epsilon = 2.58$. Second, let's compute S_{max} - the rate standard deviation:

$$S_{max} = \frac{R}{2\sqrt{3}} = \frac{1{,}000{,}000\ \ bps}{3.46} = 289{,}017\ bps$$

Then:

$$C_{stat} = \frac{1{,}000{,}000\ bps}{2} \times 100 + 2.58 \times 289{,}017\ bps \times 10$$
$$= 50\ Mbps\ + 7.5\ Mbps\ = 57.5\ Mbps$$

From this example we can see that a worst case provisioning method would require 100 Mbps of link capacity while using statistical over-subscription, a link of capacity of roughly 60 Mbps can be used to support the aggregate traffic volume with a confidence level of 99% (we are confident that 99% of the time the aggregate volume of traffic will not exceed 60 Mbps).