

Sources of Traffic Demand Variability and Use of Monte Carlo for Network Capacity Planning

Alex Gilgur (Google, Inc), Brian Eck (Google, Inc)

Abstract

When sizing any network capacity, several factors, such as Traffic, Quality of Service (QoS), and Total Cost of Ownership (TCO) are usually taken into account. Generally, it boils down to a joint minimization of cost and maximization of traffic subject to the constraints of protocol and QoS requirements. Stochastic nature of network traffic and link saturation queueing issues add uncertainty to the already complex optimization problem. In this paper, we examine the sources of traffic demand variability and dive into Monte-Carlo methodology as an efficient way for solving these problems. Other sources of uncertainty in network capacity forecasting are briefly discussed in the Attachment.

Keywords:

network; demand; capacity; link; node; monte-carlo; latency; ROC curve; QoS; simulation; throughput; traffic; concurrency; availability; node-and-link model; fast-time simulation; agent-based simulation.

“Life is amazing, isn't it? You can't ever tell what will happen. Nobody knows until they go ahead and play the game.”

– Gennifer Choldenko, *Al Capone Does My Shirts*

Introduction. Network Capacity Planning: “No Pressure?”

When planning capacity for any resource, we always have to project the future growth in demand for the resource. We may make a mistake, overestimating the demand, and then we have to deal with rightfully angry business and finance teams: physical resources start depreciating the moment they were paid for. Or we can underestimate the demand and face bounced transactions, dropped packets, and unhappy customers.

Network resources are no exception, especially in today's globally interconnected world, when we want to know who won the Olympic Gold before the champion does, and a Stanford doctor can direct a life-saving surgery from 12,000 miles across the world. The growth of network traffic over the last 10 years ensures that regardless of what we might have thought when we sized the resource, once it is in operation, we are guaranteed to sooner or later hit the ceiling. Consequently, we need to know capacity prediction intervals.

In this paper, we describe the structure of a network from capacity standpoint, discuss traffic as the main contributor to the uncertainty in capacity decisions, and describe application of the Monte-Carlo method to sizing Network Capacity. A discussion of other sources of uncertainty that leave us no choice but to “play the game” is provided as well in the Attachment.

Dealing with the Uncertainty: Monte-Carlo to the rescue!

Monte-Carlo methodology was first implemented in the 1940s during the work on the Manhattan Project. Since then, it has evolved into a general-use methodology of making statistically sound predictions,

regardless of there being an analytical solution, where computing the probability distributions for each scenario would be too complex.

In the IT world, Jeffrey Buzen blazed the trail in his groundbreaking paper in 2006 [BUZN2006] to using Monte Carlo principles for solving performance analysis problems. We are extending his work into capacity planning, having successfully applied Monte Carlo methodology for network performance analysis and capacity planning.

Definitions:

- Lead Time = amount of time between deciding to provision capacity and when that capacity is needed. Lead times are key to answering a number of supply-demand questions.
- Quality of Service (QoS) = a generic term describing the performance of a system (availability, probability of blocking, latency, error rates, jitter, etc.). Parameters are defined by what is important to the user. It is often used interchangeably with the term Service Level.
- Flow identifier = a tuple of [QoS, Source Node, and Destination Node].
- Flow = A pair of [flow-identifier, traffic]

Sizing the Network to Protect Against Demand Variation at Lead Time

Sizing the network (L3) layer of a communication network integrates protection of given physical (L1) traffic against failures, as well as topology and traffic engineering considerations.

Once the network has been sized to accommodate topology and protection against failures, we can then consider the question of protecting against demand in excess of what was predicted. This variation in traffic demand is not the within-day and within-week variation (shown in later sections), but rather, a measure of how well our forecasting engine is able to predict the target we are forecasting (e.g., weekly 95th percentile of demand).

From this standpoint, it is important to understand this variation (in our forecast versus measured demand) relative to the lead time at which we make an investment decision to add capacity. For example, if the lead time to acquire new capacity is two months, it is of interest to understand our ability to forecast traffic (and therefore the associated capacity) two months in advance.

In addition to the challenges involved in understanding the variation in traffic forecasts (see, e.g., [MAKR1998] and [ARMS2001]), it is important to note that translating this variability into network capacity forecasts adds a new dimension to the problem.

A Closer Look at an Example Node-and-Link Model Structure

A fragment of a traffic network example represented by nodes and links is shown in Figure 1. Packets traverse the nodes via links, each node and link having their intrinsic properties: retention time; capacity; throughput, reliability, probability of choosing one exit link (“degree”) over another, etc.

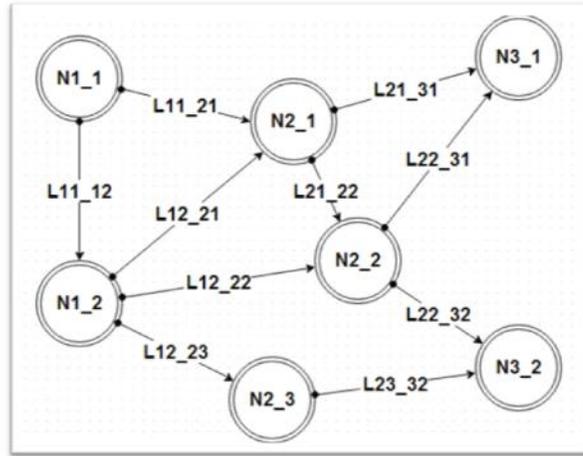


Figure 1: A fragment of a traffic network

Any node-and-link model involves activities (akin to decision making in a decision tree) on nodes, connected by links. For a data network, the nodes and the links represent the data centers, spans of fiber, routers, switches, amplifiers, etc. The protocol (control software in case of software-defined network) describes the “IF...THEN...ELSE” rules for the nodes, and each link at any moment in time is characterized by a current load, round-trip time (RTT), and a capacity.

Problem statement

We have:

1. A network
 - a. Topology, including Possible Points of Failure
 - b. Protocol(s) and Policies
2. Node A (source) to Node Z (sink) Traffic demands

We need to determine:

1. What the capacity of the links should be, in order to guarantee flow while protecting for variability inherent in the forecasting of network demand (traffic).
2. How much and how far in advance to provision capacity and how much hardware to overstock.

Additional challenge here comes from the fact that for network, traffic requirements do not directly translate into link sizes. It is therefore important to understand:

- the variation in network capacity required to support variability of traffic, and how network layer (L3) capacity requirement variability translates into physical-layer (L1) provisioning.
- How much traffic we will have on each physical link of the network.

Solutions (a brief overview)

Very often, an analytical solution exists and can be solved. A number of authors, most notably Menasce & Almeida [MENA2002], and others, have written on the subject.

However, in other cases, e.g., for a very large network, a closed-form analytical description of the network capacity would be unfeasible or inadequate. Traffic simulation is required. A number of simulation

techniques and off-the-shelf modeling software are available. They all use a similar approach of simulating the grid by introducing virtual flows parsing the network, each carrying their payload (typically number of bits per second). By sending multiple flows down the network and simulating retention times and turns / splits / losses at each node for each of them, they fill network topology to capacity, implementing optimization in some form (typically, it is the Shortest-Path-First, or SPF algorithm), and making capacity-augmentation suggestions to ensure the QoS requirements are met. An important downside of these tools and techniques is that they are deterministic and as such, cannot be directly used to completely answer the questions of variability of network capacity forecasts.

Network Traffic Predictability

Retrospective analysis is most useful when it can give important insights into how to deal with the future. It is too late when there is congestion, a lost communication, or a denial of service. We want to be able to forecast traffic on the network.

In the computing domain, traffic is usually interpreted as number of concurrent jobs, or - applying Little's Law - the product of latency and either arrival rate (making 'offered traffic') or throughput (making 'serviced traffic'). We are interested in maintaining a given level of service, defined as the lowest acceptable probability of uninterrupted operation (see, e.g., [CRFT2006]; [FERR2012(1)] and [FERR2014]). This probability is tied into concurrency and number of circuits available for serving the requests. We get congestion when we have more concurrent jobs to process than can be kept in the pipeline; therefore it is more important to know how many jobs are in the system at any given time than to know the arrival rate of the packets.

In the network domain, it is different. For a network, traffic is usually interpreted as packet traffic, measured in multiples of bits per second (Gbps), and capacity is the maximum throughput that the network or its components can carry. In most cases, we are interested in the customer getting their data fast, while the network protocols take care of ensuring delivery (description of operation of network protocols is outside the scope of this paper; [NAMB2013] is an excellent overview of the subject).

It is important to accurately predict network traffic demand. There are a number of time-series-analysis and regression techniques that can be used in traffic forecasting. A good description is given in, e.g., [MAKR1998] or [ARMS2001].

Network demand is characterized by five very important features (the first four are covered below; the fifth one would take the discussion on a tangent and is not included here; some of it is covered in, e.g., [FERR2014]):

1. It is noisy.
2. It is not stationary in time.
3. It varies by the node of origin and node of destination.
4. It varies by the QoS and latency requirements.
5. It may be bursty.

These features require advanced analysis techniques to predict variability of network demand at the hardware lead time to guide network build planning.

Noise

The high level of noise is common in any network where a wide variety of services and people are sending their communications to each other. It is especially true of the Internet traffic.

However, despite the noise, some patterns can be discerned, often following distinct cycles (annual, weekly, and daily; sometimes hourly).

Noise should not be confused with outliers - though they often make one of the most prominent contributors to noise, especially ones that are driven by big events and other predictors (big holidays; big-news events; Soccer World Cup, etc). Such outliers are not to be discarded in generation of forecasts: the network capacity planner typically will want to be able to support such events.

Because we are interested in sizing the network for high traffic, we have to agree on a percentile of the demand distribution for which we are sizing the network, as well as how to use this percentile. Typically in IT capacity planning, a forecast of the weekly 90th or 95th percentile of demand is used. Network is no exception. This imposes additional challenges, amplified if the forecasting is not done based on a statistically significant sample:

- 95th percentile of the traffic values falls on the tail of the distribution, where the highest noise is usually observed.
- The tail of any distribution is known to behave differently than its bulk. We give up information that could help us form a better judgment about data behavior.
- In the most commonly used capacity forecasting workflow, the weekly 95th percentile numbers are fed directly into the forecasting model. This imposes on the analytical engine the challenge of reconstructing the bulk from the tail or to use advanced analytical techniques to obviate the need for such reconstruction.
- Systematic patterns can hide outliers and force the higher weekly percentiles to tend towards a particular day of the week.

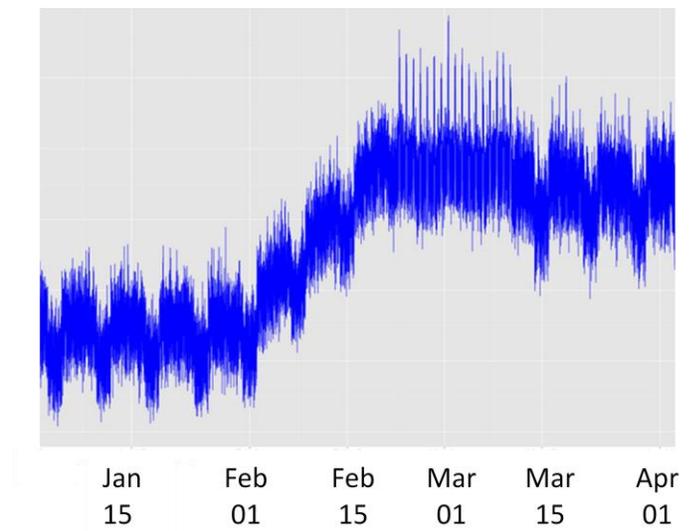


Figure 2: An illustration of data collected over the course of 3 months (the year and the vertical-axis labels are intentionally omitted to protect the Company's data).

Figure 2 illustrates the potential for losing important information: during the three weeks of 02-01 through 02-21, there was a steady growth, but if we were only using each week's 95th percentiles to evaluate the data behavior, we would have missed it and only seen step-function level changes during these weeks.

By using appropriate predictive analytics techniques in combination with domain knowledge, these challenges can be, and have been, successfully overcome; however, it is important to be aware of them.

Temporal Non-Stationarity

Network traffic is, almost by definition, non-stationary. Temporal non-stationarity is covered briefly in Attachment. An in-depth discussion involving physical-layer controls, latency, and their interactions is out of scope of this paper.

However, even if we abstract away this kind of intrinsic non-stationarity, traffic develops in patterns that are not stationary. Figure 3 is a good case in point. We see that during the week of 01-22, there was a brief increase in network usage, followed immediately by a ramp down that had a distinct weekly pattern. The ramp lasted about 3 weeks. After that came a burst in outlier activity through the end of February and most of March, followed by a level shift up with some data stabilization. This kind of temporal non-stationarity imposes an additional challenge for the analyst.

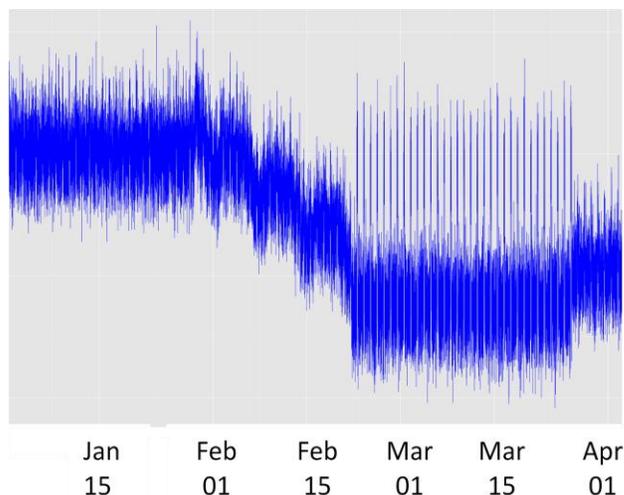


Figure 3: Temporally Non-stationary signals in network traffic

Spatial Non-Stationarity

Spatial non-stationarity of network demand manifests itself in the significant differences in the behavior of traffic as one compares different nodes among themselves.

Such differences are typically correlated with node development, data center growth and repurposing, business operations, etc. Attribution of these differences to spatial non-stationarity is further complicated by the noise and temporal non-stationarity of the flows.

An important feature of network traffic when it comes to spatial non-stationarity is flow correlations. Depending on whether the flows vary together or in the opposite directions, flow correlation may reduce or amplify spatial non-stationarity. It is important to account for it when sizing network capacity.

Importance of QoS Requirements

Different flows serve different users, and each user has his/her/its own requirements on QoS and Latency. Over time, these requirements change, as even the same user may need a different mix of traffic.

Running all these flows through the same network spans can lead to smoothing out the differences in the near term, but when projecting into the future, each flow type has to be forecasted on its own: otherwise, important growth patterns can be missed.

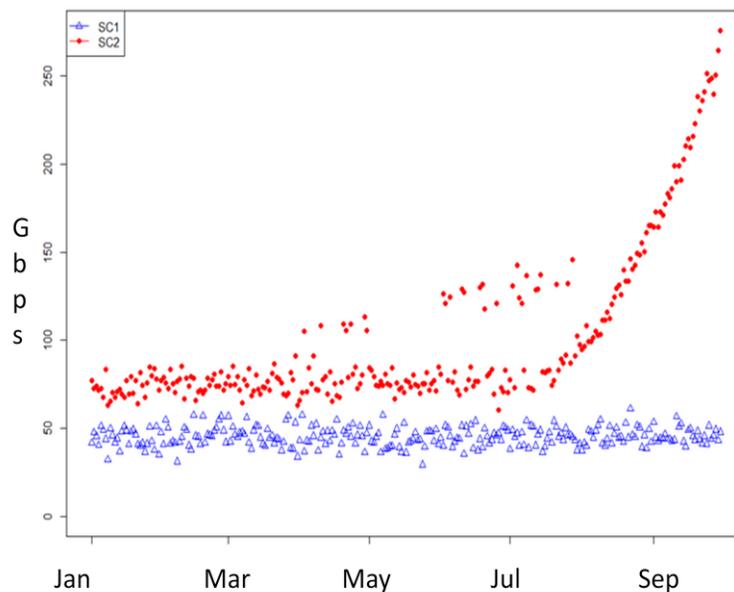


Figure 4: An illustration of the importance of capturing growth patterns in different flow separately.

Figure 4 illustrates a case where, for almost two quarters, the two flows were following a similar trajectory (one higher than the other, but parallel). In April and in June-July, 'SC1' flows started showing non-stationary behavior (a cluster of outliers), which towards the end of Q3 turned into a steep trend.

All too often forecasting is done for aggregated flows regardless of the driving force behind each of them, and as a result the corresponding network links may end up significantly undersized for such flows. We should always do the right thing by aggregating flows in forecasts based on their behavior and proceed with caution when bringing together flows with different QoS requirements. (The QoS requirements are usually dictated by business models, and their togetherness in a time series may be deceptive.) Then we can respond to the April and June clusters in SC2 by anticipating the growth we see later in the year.

Demand Forecasting

Most of the demand-forecasting challenges outlined above are surmountable by using the right statistical techniques. Any forecasting methodology produces model quality parameters and residuals whose

distribution can be used to predict the range of the forecast at any point in the future. (see, e.g., [MAKR1998], [GILG2006 & 2012], [ARMS2001]).

From the practical standpoint, we are interested in the predicted variability of network demand at the hardware lead time. It is understood that it will be different for different lead times.

Recap of what we have established so far

- Network traffic is highly variable, and we can quantify this variability.
- There is no single analytical expression translating network traffic demand into network capacity; a simulation is the practical solution, and network simulation tools exist.
- We need to know the confidence levels of the capacity predictions that we are providing.

The questions we are interested in:

- How much buffer do we really need?
- Will we have sufficient buffer built into the physical infrastructure to deliver uninterrupted services at the required service level until the new network capacity is actually built and put into operation?

In the next sections, we describe using Monte Carlo analysis to answer these questions.

Monte-Carlo simulation for a node-and-link network

The diagram in Figure 5 illustrates a probabilistic view of the network.

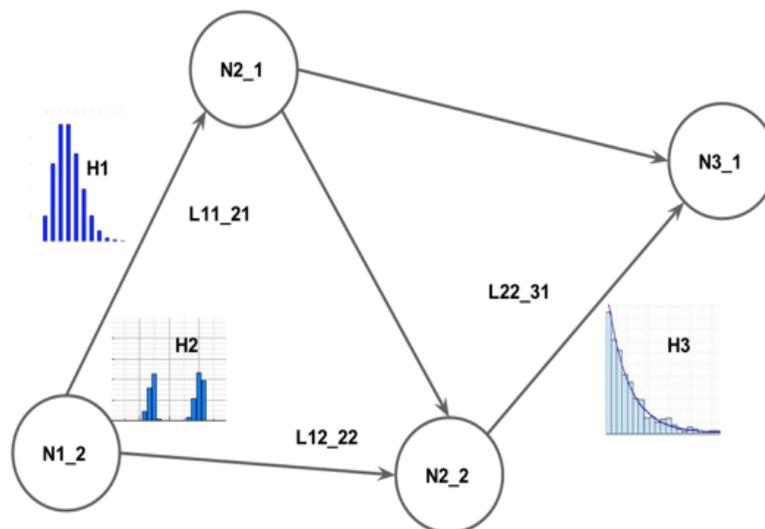


Figure 5: A fragment of the data network in Figure 1 with distributions of traffic intensities, latencies, and routing probabilities is shown for illustration.

Flows (traffic intensities) can be described as distributions like the histogram H1. However, these are physical-layer (L1) flows are derivative. They are defined by the traffic demand for the network (e.g., deliver X packets per second from Node N1_2 to Node N3_1).

Propagation and queueing delays (e.g., histogram H3) contribute to latency, while routing decisions (e.g., histogram H2) contribute to the network's ability to satisfy the demand. They both play an important role in the uncertainty of capacity forecasts, discussed in the Attachment. For this paper, their effects are assumed to be constant.

Monte-Carlo Capacity Forecasting

The uncertainties that we have discussed affect traffic predictability, making it necessary to use Monte Carlo methodology in Network Capacity planning. The Monte Carlo procedures can be very simple or very elaborate, but at the core they all involve varying the traffic demand for each flow as implied by the observations and running the capacity-sizing simulation for each value of the traffic. This leads to a set of capacity forecasts, which can then be used to directly answer the questions that were formulated in an earlier section of this paper.

A Monte-Carlo procedure for Network Capacity Planning should include the following steps:

1. Define the network topology.
2. Define the protection policy that determines the path simulated flows will take under network failure scenarios
3. Produce a "baseline" network capacity forecast for zero-error demand forecast.
4. Analyze the demand forecasting method by comparing to historical actuals:
 - a. Define lead times for network infrastructure components.
 - b. Characterize the distribution of forecast errors (comparing forecast at lead time to values measured during the same calendar time frame) for each flow in the network.
 - c. Use this distribution to generate N pseudo-random errors for each flow. Standard inverse-cdf or accept/reject methods can be used. For each flow, each sample is a single value of the error drawn from a distribution representative of the errors characterized from historical data.
 - d. It is possible that there may not be enough samples for each flow to generate the pseudo-random errors for the Monte-Carlo simulation. Special nonparametric aggregation techniques can be used in this case, aggregating the QoS of all flows for each lead time into groups whose forecast biases are similarly distributed. Each flow identifier will then be characterized by a distribution representative of the actual data, grouping together flows with similar forecastability.
5. Fix a point in time as the current forecast date (most recent forecast made), and for the fixed lead time decided upon in step 4, apply the pseudo-random errors to the baseline forecast, to generate N samples of traffic values (for the time period of interest).
 - a. This generates N vectors of (flow-identifier, traffic) pairs.
6. Use the same network simulation engine as was used for baseline to generate the corresponding capacity required to support each sample's traffic for given topology and failure protection policy.
 - a. This will produce the capacity such that traffic forecast (including its randomly-generated errors) on each link is supported at its required QoS.
 - b. Given sufficient number of samples N, and provided that the same simulation engine is used for all samples and for the baseline, flow interactions will be accounted for.
7. The resulting set of N capacity forecasts has variation as implied by traffic forecast at lead time.
8. These capacity forecasts can be used as an input to standard supply chain techniques to optimize QoS and TCO and size hardware buffers for optimal levels of service.

Implementation

We have recently completed evaluation of ranges of a large WAN capacity sizing process using the Monte-Carlo approach that we present in this paper. Without going into proprietary details and publishing proprietary data, it can be said that the methodology works, and we were able to successfully quantify the ranges of capacity requirements demanded by the uncertainty in traffic prediction, as well as formulate the specifications for each flow's traffic prediction intervals.

One of the key observations from this study was that even if demand variability may be very large, and consequently the prediction interval of demand forecast may be very wide, the network simulation dampened that variability, providing forecasts that were for the most part within tolerance imposed by the network build requirements. It is not an unexpected result: a network of N links potentially can have up to $N*(N-1)$ nodes, meaning that we are putting up to $N*(N-1)$ flows through N bidirectional links. An inevitable flow multiplexing will occur, smoothing out the variability of the flows.

It is granted that different networks and different simulators behave differently, and their sensitivity to traffic predictability may be, and will be, different as well. Monte-Carlo approach can be used in such cases to compute their use-case-specific tolerances for demand forecast variability.

Conclusions

An in-depth analysis of network traffic variability and its effect on the uncertainty of network capacity forecasts has been presented. Other sources of uncertainty are covered in the Attachment. Monte-Carlo methodology has been described, and its use for network capacity planning has been discussed. A way of including it in network simulations has been outlined.

Acknowledgements

Authors would like to express their gratitude to C. Stephen Gunn, Harpreet Chadha, Rajesh Krishnaswamy, and Thomas Olavson for reviewing this paper prior to publication.

Attachment: Other sources of variability in Network Capacity

This attachment provides a brief discussion of other sources of variability in network capacity.

Latency

Latency is one of the key variables in network performance.

Latency of a network segment can be interpreted as the sum of speed-of-light latencies of the links along the path and the queueing latencies of the sending and receiving devices on the network. An excellent recent overview of data-transfer latencies across a network is presented in the CMG'13 proceedings [VILL2013].

Speed-of-light latency, often going into the calculation of network link's IGP metric for data networks, which in turn is used either as the cost function or a constraint for network flow optimization, is the easiest to understand component of latency: it is directly proportional to the length of the link and therefore can be measured as the intercept on the latency-load performance curve (Figure A1).

Other components of latency are less traceable, because their behavior, including signal processing time, is strongly correlated with the load, number of parallel sending and receiving ports, the router's performance

characteristics, and the load balancing method employed among the routers, as well as within each router. For more details, see, e.g., [VILL2013], [GILG2013], [GUNT2009], and others.

Network capacity numbers are typically used as the signal for sizing the L1 (physical) layers of the network, based on which network components (fiber, routers, amplifiers, etc.) can be sized so as to keep the packet arrival rate below the knee (Figure A1), thus obviating the need to account for latency.

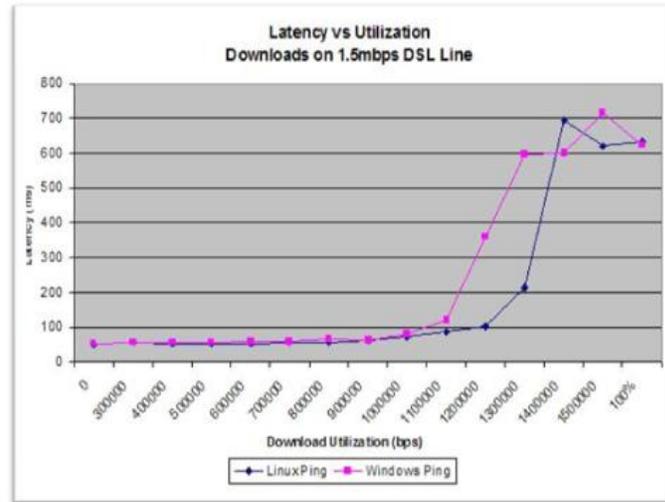


Figure A1: A Performance Characteristic (latency-load ROC) illustration. Here latency (vertical axis) is measured as the average time to download a file at a given bit rate (horizontal axis) across a DSL line.

However, being able to predict latency given load in capacity planning would allow the network simulation engine to distribute simulated flows in a manner that more accurately resembles the actual flow distribution in a physical network. A regression for latency-load ROC curve (Figure A1) would be very helpful in that.

The latency-load performance (ROC) curve (Figure A1) is well described by a hyperbola (see, e.g., [GUNT2009], [FERR2014]). However, fitting a hyperbolic curve into load-latency data based only on the low-load numbers is a daunting task even without noise, making it difficult to compute parameters of such regression.

Because of these difficulties, off-the-shelf network simulation tools account for load-driven latency increase in their simulations, but with a lot of caveats, from limitation of network loading (e.g., no more than 80%) to assumptions of distributions that may or may not be easily substantiated (e.g., saying that an adjacency can be approximated by a G/M/m queue). It is a field waiting to be explored; in this paper, we focus on traffic as the primary source of variability.

Network Protocols and Topology

Network protocols are rule-based and therefore deterministic by nature. However, like all large dynamical systems dealing with discrete signals, networks can, and do, exhibit stochastic behavior. Traffic variability adds to the uncertainty of network behavior. We need to be aware of this interaction, as it adds to the complexity in analyzing historical behavior of traffic on networks.

Network topology too plays an important role in the stochastic nature of the traffic: there are multiple paths on the network to deliver a packet from point A to point Z, potentially leading to unpredictable behaviors, where traffic may swell up in one location while running “dry” in another.

Software defined networks (SDNs) provide a better control of the stochastic component of traffic than traditional networks do; however, to properly provision an SDN, one needs to know the algorithms used in the SDN control software, which can be proprietary.

Physical Availability of L1 and L2 components

This category includes link availability, reliability of communication network hardware components, as well as their protection level against accidents such as fiber cuts and natural disasters. The standard methods from reliability theory apply here. The risk level is computed as the product of the probability of an event (failure) happening and the magnitude of its business impact.

Availability is an independent variable in the network capacity model and can easily be isolated for independent analysis. Its interaction with traffic and effect on network capacity is an interesting field that, for a network, cannot be resolved analytically. Network component availability, like all problems in the reliability domain, presents another excellent use case for Monte-Carlo analysis.

References

1. [MENA2002] Capacity Planning for Web Services: Metrics, Models, and Method. Daniel A. Menasce; Virgilio A.F. Almeida (2002). Prentice Hall, 2002.
2. [FERR2012(1)] A Note on Knee Detection. Ferrandiz, Josep and Gilgur, Alex. Las Vegas : CMG 2012
3. [CRFT2006] Utilization is Virtually Useless as a Metric! Cockcroft, Adrian. International Conference of the Computer Measurement Group (CMG'06). Reno, NV, 2006.
4. [GILG2013] Little's Law assumptions: “But I still wanna use it!” The Goldilocks solution to sizing the system for non-steady-state dynamics. Alex Gilgur (Online) MeasureIT Issue 100, June 2013
5. [FERR2012(2)] Level of Service Based Capacity Planning. Ferrandiz, Josep and Gilgur, Alex. International Conference of Computer Measurement Group (CMG '12). Las Vegas, NV, 2012.
6. [FERR2014] Capacity Planning for QoS. Ferrandiz, Josep and Gilgur, Alex. Journal of Computer Resource Management. Issue 135 (Winter 2014). pp. 15-24
7. [BUZN2006] New Perspective on Benchmarking, Modeling and Monte Carlo Simulation: Operational Analysis 2.0. Jeffrey P. Buzen. International Conference of Computer Measurement Group (CMG'06). Reno, NV. 2006.
8. [GUNT2009] Mind Your Knees and Queues - Gunther, Neil (2009). MeasureIT, Issue 62, 2009
9. [MAKR1998] Forecasting: Methods and Applications. Makridakis, Wheelright, Hyndman. Wiley, 1998.
10. [GILG2012] Time-Series Analysis: Forecasting + Regression: And or Or? A.Gilgur, J.Ferrandiz, M.Beason (2012) International Conference of the Computer Measurement Group (CMG'12), Las Vegas, NV.
11. [GILG2006] A Priori Evaluation of Data and Selection of Forecasting Model. A.Gilgur, M.Perka, W.Fuller (2006) International Conference of the Computer Measurement Group (CMG'06), 2006, Reno, NV.
12. [ARMS2001] Principles of Forecasting: A Handbook for Researchers and Practitioners. J. Scott Armstrong (2001) Springer, 2001.
13. [NAMB2013] Network Performance Engineering (2013). Manoj K. Nambiar. 39th International Conference of the Comptuer Measurement Group (CMG'13) La Jolla, CA.
14. [VILL2013] Performance Evaluation of Big Data Transmission Models (2013). Adam H. Villa, Elizabeth Varki. 39th International Conference of the Comptuer Measurement Group (CMG'13) La Jolla, CA.