Proceedings of the 2013 Winter Simulation Conference
R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, eds

PROJECTING NETWORK LOADING OF CORRELATED TRAFFIC STREAMS UNDER HIGH GROWTH

Timothy Wetzel
Timothy Lortz
Ashleigh Thompson

Booz Allen Hamilton
308 Sentinel Drive
Annapolis Junction, MD 20701, USA

ABSTRACT

The optimal design of information technology systems is a multi-objective problem requiring accurate prediction of both normal and peak loading conditions. The calculation of future peak loading requires the projection of average loading as well as a detailed understanding of parameter interactions. The failure to properly account for cross-source correlations for many component systems can lead to a significant underestimation of the total system variability. We present an IT network model that projects both average and peak loading for a system that is expanding traffic at multiple data sources. The average loading is modeled by independently calculating the traffic growth for all data sources. The future cumulative load distribution is calculated by taking into account historical sources of parameter correlations and scaling to account for future average growth. Finally, the duration of the peak states is calculated by simulating a vector Fractional ARIMA time series model.

1 INTRODUCTION

Modeling and simulation is often used to assist in the optimal purchasing of hardware and the design of network architecture. The performance of IT systems is dependent on multiple factors which include the average and peak loading of critical components in the network. The average load affects both the hardware necessary for data storage and also determining the optimal architecture of the physical data network. The magnitude and duration of peak utilization periods is required for determining the necessary I/O for network components as well as the sizing of data buffers needed to prevent data loss and maintain Quality of Service.

The network problem presented here projects the average and peak loads for a network that is going to expand front end traffic demands by a factor of four. The current average volumes are parameterized using historical data for the entire end-to-end process flow; data enters the network, is transformed by network components which can spawn child flows, and then is finally stored in a data repository. The projection of future volumes as a function of traffic increases is presented using multiple methodologies that include both linear scaling and also sublinear power law fits in order to demonstrate likely ranges of future network behavior.

2 PROJECTING CUMULATIVE DISTRIBUTION FUNCTION OF TRAFFIC LOADING

The modeling of peak loading requires knowledge of both the future average loads and also the current variability as a function of average load. The projection of the total daily network variability requires the modeling of multiple layers of parameter interactions. The future traffic intensities at an existing source are strongly correlated with the current traffic at that source. These correlations causes the future variance of a single source to scale quadratically with any growth in the average traffic volumes.

In addition to future data traffic of a source being strongly correlated with the existing traffic, there is also a weak, but positive, cross-correlation between the traffic at different sources across the network. The effect of cross-source correlations on the total system variability is dependent on the number of different data sources and for the network presented here, the cross-correlations increase the standard deviation of total network load by approximately 1.5. The correlations are primarily caused by day-of-week effects with most sites having high network utilization during the week and low data volumes during the weekend. We modeled the day of week effect of traffic intensities by using appropriately scaled bimodal distributions.

The inclusion of both correlations within a source and cross-source covariance improved the projection of peak by a factor of 3. The using a bimodal distribution to separate the effects of weekdays and weekends improved the modeling of system variability by reducing the width of the confidence interval for the standard deviation by a factor of 2.

The cross-data source covariance accurately models the entire probability distribution function but fails to represent the correct day-to-day variability. Independently sampling from the probability distribution function significantly overestimates the variability and thus underestimates the time duration of peaks. To accurately model the day-to-day variability, we use Fractional ARIMA time series models to capture the trends for the autocorrelation function of the total network load. The Fractional ARIMA model used in this study contained a one-day autoregressive term to capture short term effects, a seven-day seasonality term to capture weekly affects and an infinite, and a fractal differencing parameter that captured the long-range patterns. The combination of these terms yields a stationary time series with a slowly decaying autocorrelation function that has local maximums at modular 7. The use of a Fractional ARIMA model for data networks is consistent with observations from physical phenomena like hydrology and forestry (Hosking 1984), as well in arrival times in packet switched networks (Leland et al. 1994).

The traffic loads of the data sources were modeled individually and the total load was calculated by summing the individual time series. The use of a self-similar time series approach significantly improved the day-to-day variability and was superior to random sampling of the daily probability distribution function. Both the Monte Carlo sampling and the Fractional ARIMA projected the same number of days in a peak state, but they differed in the distribution of the clustering of peak days. Random sampling of the probability distribution function projected that almost all peak events would last only a single day whereas the Fractional ARIMA model accurately projected the distribution of peak states lasting one to five days.

REFERENCES

Hosking, J.R.M. 1984. "Modeling Persistence in Hydrological Time Series Using Fractional Differencing". *Water Resources Research* 20(12):1898-1908.

Leland, W.E., M. Taqqu, W. Willinger, D.V. Wilson. 1994. "On The Self-Similar Nature of Ethernet Traffic (Extended Version)". *IEEE/ACM Transactions on Networking* 2(1):1-15.