Predicting TCP Throughput Using Time Series Forecasting

Yinan Xue, Miguel Rio

Dept. of Electronic & Electrical Engineering, University College London

Abstract: Forecasting end-to-end throughput accurately can be vital for several applications. Peer-to-peer applications like BitTorrent or Skype will be able to select the peers more efficiently if the throughput in the short-term is known. On the other hand applications can adapt their behaviour accordingly, if they know how the end-to-end throughput is going to vary in the lifetime of a connection. This paper focuses on predicting throughput using time series forecasting techniques: Exponentially Weighted Moving Average and Holt-Winter. In this work, current data of throughput were collected using the Planet-lab simulator. It was found that for both EWMA and Holt-Winters, the accuracy patterns vary a lot from path to path. After comparing the prediction errors of EWMA and Holt-Winter, we demonstrated that if the data shows trend and seasonality, the prediction error of Holt-Winter is much lower than that of EWMA, while in other cases the accuracy of EWMA and Holt-Winter are quite similar. Finally, we analysed the correlation between the prediction and RTT. We identified that, for the predictions obtained from both EWMA and Holt-Winter, the correlation follows a power function.

1. Introduction

Throughput forecasting is very important and widely used in many applications now. For example, for server selection and overlay route selection, a throughput prediction is needed before the flow starts [1,2]. It is because rerouting an established TCP connection to a different network path or server can cause problems such as migration delays, packet reordering, and re-initialization of the congestion window.

With all this motivation, lots of work has been done in this field. Related work includes making predictions according to the network characteristics [3, 4, 5] and developing a new integrated frame work to predict the bandwidth of a TCP flow [6]. Recently, Q. He, C. Dovrolis, M. Ammar have classified TCP throughput prediction techniques into two categories: Formula-Based (FB) and History-Based (HB) and provide insight regarding the factors that affect the predictability of large transfer TCP throughput in a given path [7].

This paper focuses on predicting the throughput of a bulk TCP transfer using time series forecasting techniques: Exponentially Weighted Moving Average (EWMA) and Holt-Winters. The work consists of two main tasks: Data collection and Algorithm experimentation. The data collection was executed using the Planet-lab simulator [8]. This allows collecting a significant amount of data on end-to-end throughput between any two pairs of nodes in the simulator. After the data was collected, we tested the time series forecast of EWMA and Holt-Winters and evaluated their accuracy under different conditions. Finally, we correlated the average throughput with the RTT collected using ping and traceroute.

2. Data Collection

In the project, PlanetLab is used as the simulator for data collection. PlanetLab is an open platform for developing, deploying, and accessing planetary-scale services. It is a collection of machines distributed over the globe. All of the machines are connected to the Internet. Research groups are able to request a PlanetLab slice in which they can experiment with a variety of planetary-scale services. The advantage of using PlanetLab is that the experiment is performed with new services under real-world conditions, and at large scale. All PlanetLab machines run a common operating system. This allows an application to run across all (or some) of the machines distributed over the globe, where at any given time, multiple applications may be running in different slices of PlanetLab.

The data was collected mainly on 8 paths that interconnect 16 nodes of planet-lab. The nodes used were located in US, UK, Brazil, Germany, France, Japan, and China. The 8 paths are: (1) Campina Grande, Brazil-Fife, Scotland (2) Paris, France-Fife, Scotland (3) Karlsruhe, Germany-Manchester,UK (4) Tokyo, Japan- Manchester, UK (5) Beijing, China – Lancaster, UK (6)Columbia US-Manchester,UK (7)Berkeley, US – London, UK (8) Pennsylvania, US- Lancaster, UK. Iperf [9] was installed in every node to measure the current throughput. The data were collected for a five-day time interval.

3. Time Series Forecasting

The Time Series Forecasting is based on the current value of the throughput over a specified time period.

It is possible where large TCP transfers are performed repeatedly over the same path [7]. In this part, Exponentially Weighted Moving Average and Holt-Winters are mainly used for the prediction. We will assume that these prediction algorithms consider a fixed window size for the recent traffic history.

3.1 Exponentially Weighted Moving Average (EWMA)

3.1.1 EWMA Predictor

An exponentially weighted moving average is an exponentially weighted mean of previous data points.

$$\hat{X}_{i+1} = \alpha X_i + (1-\alpha)\hat{X}_i \quad (1)$$

The parameter α is the weight of the last measurement. A higher α cannot smooth out the measurement noise whereas a lower value is slow in adapting to changes in the time series [9].

3.1.2 Results

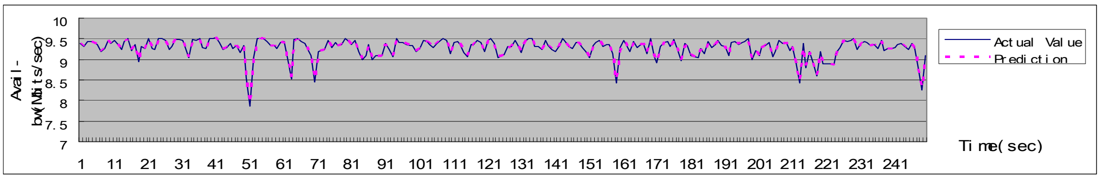

Figure 1: Comparison of actual values with the predictions (France-UK)

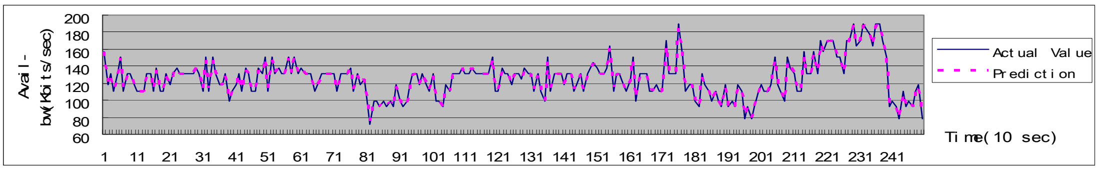

Figure 2: Comparison of actual values with the predictions (Brazil-UK)

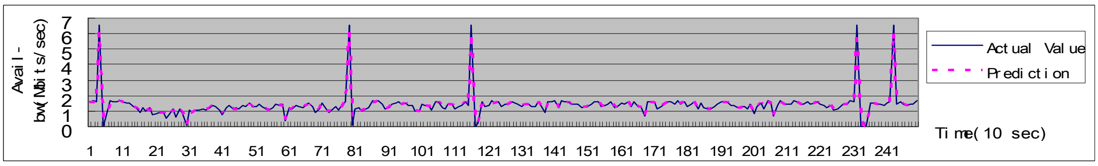

Figure 3: Comparison of actual values with the predictions (Berkeley, US-UK)

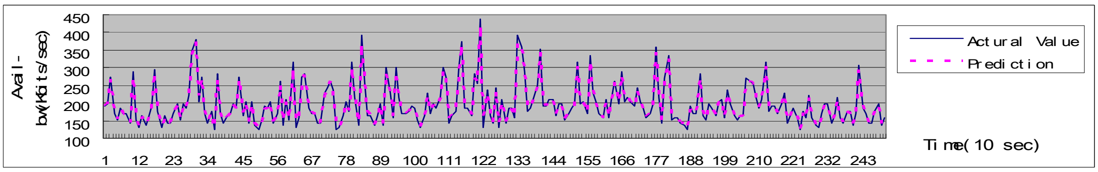

Figure 4: Comparison of actual values with the predictions (China-UK)

Figure 1-4 are the comparisons of actual values with the predictions of four paths. The results showed that the prediction errors of EWMA are quite low. We can also note that the accuracy of EWMA varies a lot from path to path.

3.2 Holt-Winters

3.2.1 Holt-Winters Predictor

(1) The basic equations for their method are given by [10]:

$$L_t = \alpha(Y_t/S_{t-c}) + (1-\alpha)(L_{t-1}+T_{t-1}) \quad (2) \qquad T_t = \beta(L_t - L_{t-1}) + (1-\beta)T_{t-1} \quad (3) \qquad S_t = \gamma(Y_t/L_t) + (1-\gamma)S_{t-c} \quad (4)$$

$$F_t = (L_t + mT_t)S_{t-c+m} \quad (5)$$

where y is the observation, L is the smoothed observation, T is the trend factor , S is the seasonal index , F is the forecast at m periods ahead , t is an index denoting a time period , and α, β, and γ are constants that must be estimated in such a way that the MSE of the error is minimized.

(2) Initial Values for the Data Level

$$L_0 = \frac{Y_1 + Y_2 + \dots + Y_c}{c} \quad (6)$$

This value of the Data Level is an average of the observations of the first season.

Initial Values for the trend factor

$$T_0 = \frac{Y_{c+1} + Y_{c+2} + \ldots + Y_{c+c} - Y_1 - Y_2 - \ldots - Y_c}{c^2} \quad (7)$$

Initial Values for the Seasonal indices

For k=1,2,3,…,c-1,c

$$S_{k-c} = \sum_{i=0}^{1} \frac{Y_{k+i \bullet c}}{\sum_{j=k+i \bullet c}^{(i+1) \bullet c} Y_j} \quad (8)$$

This value of the trend factor is mainly an average of the differences in the observations of the first two seasons divided by the length of 1 season.

This value of the seasonal index is basically making an average of the observed values of every season and then making an average of the period n of every season divided by the average of the observed values for its season.

3.1.2 Results

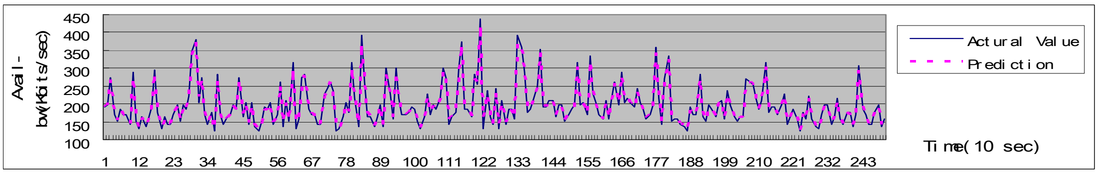

Figure 5: Comparison of actual values with the predictions (France-UK)

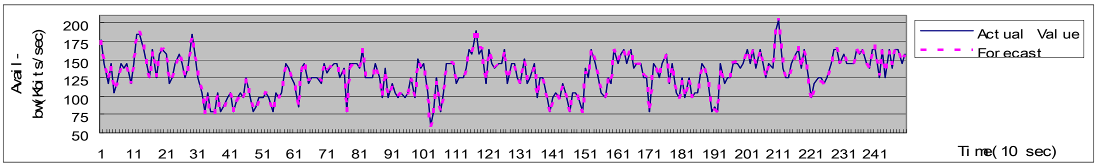

Figure 6: Comparison of actual values with the predictions (Brazil-UK)

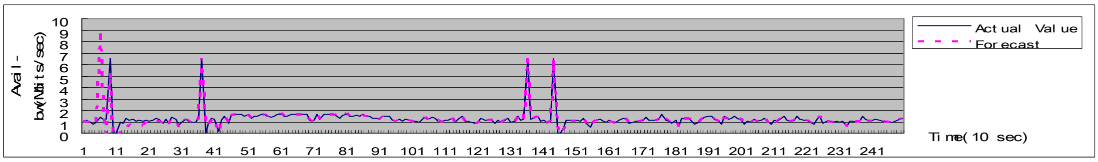

Figure 7: Comparison of actual values with the predictions (Berkeley, US-UK)

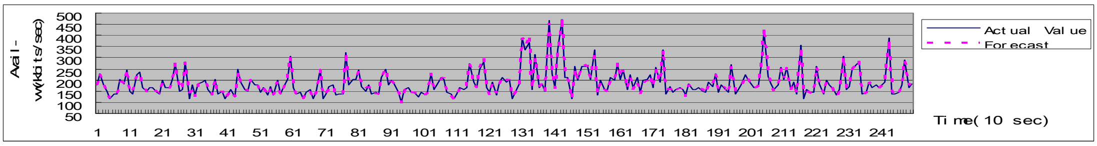

Figure 8: Comparison of actual values with the predictions (China-UK)

Figure 5-8 are the comparison of actual values with the predictions of four paths. The results showed that Holt-Winters works if the data show trend and seasonality. Different paths have absolutely different patterns of prediction accuracy.

3.3 Error Comparison of EWMA with Holt-Winters

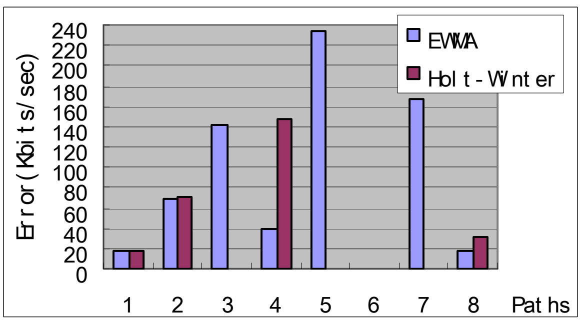

Figure9: Error Comparison of EWMA with Holt-Winters

Figure 9 is the comparison of prediction errors calculated using MSE.The chart shows that if the data show trend and seasonality, the prediction error of Holt-Winters is much lower than that of EWMA, while in other cases the accuracy of EWMA and Holt-Winters are quite similar. We can also find from Figure 9 that, for both EWMA and Holt-Winters, the accuracy patterns vary a lot from path to path.

4. Correlation of Available-bandwidth and RTT

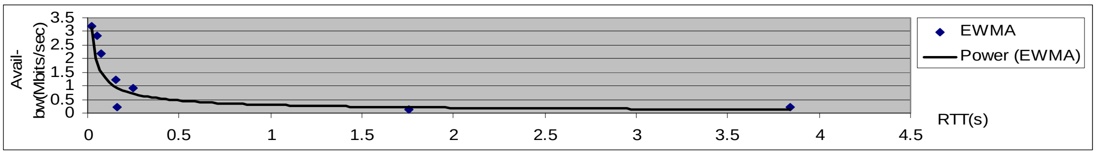

Figure 10: Correlation between EWMA Prediction and RTT

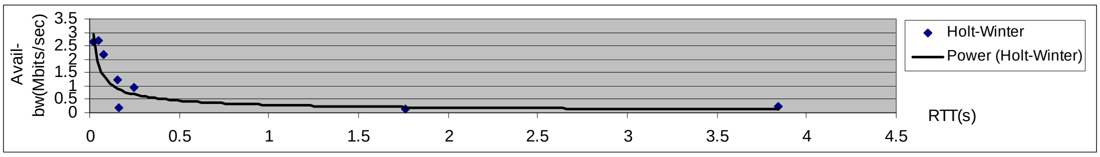

Figure 11: Correlation between Holt-Winters Prediction and RTT

Figure 10 shows the correlation between EWMA prediction and RTT. Figure 11 shows the correlation between Holt-Winters prediction and RTT. It is found that both of them follow a power function.

5. Conclusion

This paper collected the data of current throughput and made predictions using time series predictor: EWMA and Holt-Winters. After comparing errors of predictions, we identified that if the data shows trend and seasonality, the prediction error of Holt-Winters is much lower than that of EWMA, while in other cases the accuracy of EWMA and Holt-Winters are quite similar. It is also demonstrated that for both EWMA and Holt-Winters, the accuracy patterns vary a lot from path to path. We also evaluate the correlation between the prediction and RTT and found that it follows a power function.

Throughput forecasting is the subject of ongoing research and development. The final goal may be to develop algorithms, and corresponding software that enables any application to obtain an estimation of the end-to-end throughput and quality of service with the minimal amount possible of information about the node it wants to connect to. This will likely involve the definition of a protocol of exchange of information between nodes that will share the data collected by a particular node with nodes close to it.

6. Acknowledgments

I would like to thank my supervisor Miguel Rio for his patient and valuable instructions, Zhaohong Lai for help with UNIX, Kerry Jean for the help with the data collection.

5. Reference

[1] Dongyu Qiu and R. Srikant. Modeling and Performance Analysis of BitTorrent-Like Peer-to-Peer Networks. http://www.cse.cuhk.edu.hk/~cslui/STUDY_GROUP/BT_MODEL.pdf.
[2] Mort Naraghi-Pour. Bandwidth and Bu_er Dimensioning for Guaranteed Quality of Service in Wireless ATM Networks. http://www.ee.lsu.edu/mort/publications/wcnc2_00.pdf.
[3]_Allen B. Downey: An empirical model of TCP performance. In the 13th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2005). Atlanta, GA, USA, 27-29 September 2005.
[4] J. Padhye, V.Firoiu, D.Towsley, and J. Kurose. Modeling TCP Throughput: A Simple Model and its Empirical Validation. IEEE/ACM Transactions on Networking, 8(2):133–145, 2000.
[5] Tsung-i Huang, Jaspal Subhlok. Fast pattern-based throughput prediction for TCP bulk transfers. In the 5th International Symposium on Cluster Computing and the Grid (CCGrid 2005). Cardiff, UK, 9-12 May, 2005.
[6] S. Vazhkudai, J.M. Schopf, and I. Foster, Predicting the Performance of Wide Area Data-Transfers, Proceedings of the 16th Int'l Parallel and Distributed Processing Symposium (IPDPS 2002), to appear.
[7]Q. He, C. Dovrolis, M. Ammar. On the Predictability of Large Transfer TCP Throughput. In the Proc. of ACM SIGCOMM. Philadelphia PA, August 2005.
[8] http://www.planet-lab.org/
[9] http://dast.nlanr.net/Projects/Iperf/
[10]Maria Papadopouli, Elias Raftopoulos, Haipeng Shen. **E**valuation of short-term traffic forecasting algorithms in wireless networks**.** 2nd Conference on Next Generation Internet Design and Engineering, Valencia, Spain, April 3-5, 2006.
[11] Sang C. Suh, Sam I. Saffer, Dan Li, Jingmiao Gao. A new insight into prediction modeling system. In the Integrated Design and Process Technology, IDPT-2003.